

Agent-Based Control for Decentralised Demand Side Management in the Smart Grid

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

ABSTRACT

Central to the vision of the smart grid is the deployment of smart meters that will allow autonomous software agents, representing the consumers, to optimise their use of devices and heating in the smart home while interacting with the grid. However, without some form of coordination, the population of agents may end up with overly-homogeneous optimised consumption patterns that may generate significant peaks in demand in the grid. These peaks, in turn, reduce the efficiency of the overall system, increase carbon emissions, and may even, in the worst case, cause blackouts. Hence, in this paper, we introduce a novel model of a Decentralised Demand Side Management (DDSM) mechanism that allows agents, by adapting the deferment of their loads based on grid prices, to coordinate in a decentralised manner. Specifically, using average UK consumption profiles for 26M homes, we demonstrate that, through an emergent coordination of the agents, the peak demand of domestic consumers in the grid can be reduced by up to 17% and carbon emissions by up to 6%. We also show that our DDSM mechanism is robust to the increasing electrification of heating in UK homes (i.e., it exhibits a similar efficiency).

Categories and Subject Descriptors

I.2.11 [Computing Methodologies]: Artificial Intelligence—*Distributed Artificial Intelligence*

General Terms

Agents, Multi-Agent Systems

Keywords

Energy, Demand-side Management Electricity, Multi-Agent Systems, Agent-Based Control, Agents.

1. INTRODUCTION

The creation of the Smart Grid has been posed as one of the greatest challenges of this century, as countries face dwindling non-renewable energy sources and the adverse effects **Cite as:** Agent-Based Control for Decentralised Demand Side Management in the Smart Grid, S. D. Ramchurn, P. Vytelingum, A. Rogers, N. R. Jennings, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems – Innovative Applications Track (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 5-12.

Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

of climate change due to carbon emissions [12]. The vision of a Smart Grid includes technologies that enable the efficient integration of intermittent renewable energy sources (such as wind or solar energy) and electric vehicles, and will reduce demand by allowing consumers to better manage how electricity is used, stored, and delivered. One of the key underpinnings of this endeavour is the concept of the *smart meter* which aims to manage the devices in the home to minimise inefficiencies in usage and maximise the user's savings. Smart meters also aim to interact with the grid in order to help reduce peaks in demand (which would otherwise lead to instability in the grid, higher energy costs and higher carbon emissions) and keep up with variable output from wind or solar energy generators [1]. If these ideal features of smart metering materialise, they may lead to significant reductions in energy costs and carbon emissions while guaranteeing security of supply.

To date, smart meters have mainly been designed to act as information provisioning devices that tend to leave it to the user to manage devices in the home, with the hope they will reduce their energy demands. In addition to this, Demand Side Management (DSM) technologies have also been developed to alter the behaviour of users by either charging them for using electricity at peak hours (potentially leading to other peaks at cheaper hours) or by inducing the devices, via the smart meter, to turn off (or on) when the network signals them to (either through price signal or the frequency of AC power) [6, 7].¹ While such DSM techniques have been shown to bring about significant improvements on a small number of houses, it is unclear how such technologies will scale when smart meters are rolled out to millions of homes or buildings nationwide.² In particular, the centralised management of even thousands of smart meters is likely to be a complex task that may require intruding upon users' privacy to cater for all homes, each with its own specific set of devices, usage profile, and preferences set by the owner (e.g., thermostat settings or times at which she wishes to switch on energy intensive devices such as washing machines). Moreover, as we show in this paper, simply leaving smart meters to react to price or frequency fluctuations can cause all devices to respond *at the same time* (e.g., every user switching

¹A change in frequency is caused by unmatched supply and demand in the grid. Maintaining the frequency of electricity (50Hz in UK) is important because devices running on AC power such as electric motors in household appliances are optimised for this frequency and may be damaged if it deviates too far.

²As per the Climate Change Act of 2008, the UK Government has committed to installing smart meters in all of the 26M homes in the UK by 2020 [2].

on the air conditioning unit or washing machine at the same time), thus generating peaks in demand that impact on the system. Finally, the fact that increasingly more and more features of the home are likely to be electrified in the future [3] (e.g., the use of heat pumps for space and water heating), means that more significant peaks may be created due to the reactive behaviour of the smart meters. Thus, unless appropriate control strategies are implemented to allow smart meters to *coordinate*, they may well generate greater peaks in demand than before, leading to a more stressed grid and, in the worst case, blackouts.

Against this background, in this paper, we develop and evaluate a model of *decentralised demand side management* (DDSM) to coordinate large populations of autonomous agents representing individual smart meters. In more detail, we model the *smart home* as being composed of a number of deferrable loads (controlled by an agent that optimises the deferment of these loads so as to maximise the comfort in the home and minimise energy costs) as well as non-deferable loads. Moreover, we design mechanisms for agents to *adapt* (instead of reacting as in traditional DSM mechanisms) the deferment of their loads. Thus, this paper advances the state of the art in the following ways. First, we provide a model of the smart home where the deferment of loads is optimised using mathematical programming techniques. Second, we provide motivating examples that illustrate the key requirements to implement fully decentralised agent-based control strategies and thereon show that our model of DDSM meets these requirements. Third, we empirically evaluate our DDSM mechanism (on a population of 5000 agents) and show that using DDSM leads to the *emergent coordination* of agents (without directly communicating with each other). Based on the average load profile of 26M UK homes, the agents converge to an equilibrium where the peak demand of the domestic consumers is flattened by up to 17% and carbon emissions are reduced by up to 6%. Furthermore, using evolutionary game theoretic techniques, we also show that it is always profitable (i.e., it is a Nash Equilibrium) to allow agents to control the deferment of devices as the proportion of smart homes in the population increases. Finally, we demonstrate that our DDSM mechanism would remain effective even in a future with significantly more electrified heating.

The rest of this paper is structured as follows. In Section 2, we review the literature and in Section 3, we describe our model of the smart home. Section 4 presents the mathematical programming solutions we propose to the optimisation problems that may need to be solved by the smart meter. We present the requirements for implementing the DDSM and then, our DDSM in Section 5. In Section 6, we empirically evaluate the performance of the DDSM in terms of its impact on the smart home and the grid performance. Section 7 concludes.

2. BACKGROUND

Scheppe *et al.* introduced the concept of DDSM and were the first to propose a design for the smart meter which included advanced functions to optimise the schedule of loads, the prediction of demand in the home, and the prediction of weather conditions among others [10].³ While being far ahead of their time, their work mainly dealt with predicting

³Prior to this, they had provided techniques to (i) control demand and supply in the network through the use of spot pricing of electricity [11] (ii) perform demand side manage-

the system behaviour, given the implementation of smart meters, using closed form solutions that required approximating and homogenising the behaviours of actors involved in the system. In contrast, in this paper, we present an agent-based approach that allows us to model each individual home in the system and to simulate large numbers of such homes (potentially in the order of millions given enough computational resources) in order to analyse system performance. In so doing, we are able to establish the incentives of the users to adopt smart metering technology using agent-based simulations and evolutionary game theoretic (EGT) techniques.

Similar to our work is that of Vytelingum *et al.* [13] who presented a model of smart meters controlling storage devices in the home and showed how doing so improves the performance of the system at large. However, they do not deal with more complex deferrable loads and managing the comfort in the home, which are key issues in DSM. Hence, we believe our paper generalises their approach by formalising and evaluating the control mechanism needed to elicit good equilibria in the system when performing demand side management.

In addition to these theoretical studies, recent DSM trials by the GridWise Project⁴ have shown that market-based control techniques, where homes respond to real-time pricing (RTP), can reduce peak demand (when prices are high) to some degree [6, 7]. Theoretical studies of this setting also predict similar benefits [8]. However, these approaches simplify the behaviour of the devices in the home to *reactively* respond to prices rather than predicting and planning use. This uncoordinated behaviour, in turn, leads to peaks in demand being shifted to periods when the prices are low. In contrast, in our DDSM, the agents *adapt* the behaviour of the devices which allows them to coordinate, in an emergent fashion, to flatten demand.

3. THE SMART HOME MODEL

In this section we present a model of an agent that optimises the home’s energy usage by managing loads to maximise savings while mitigating the impact on the user’s lifestyle (or comfort). Now, it is important to note that loads can be classified into two categories: those that can and those that cannot be deferred to later times. Examples of the latter include lighting, entertainment devices, phone charging and computer usage and of the former, include washing machines, dishwashers, boilers and fridges. Hence, the agent is only able to control some of the devices in the home without impacting too much on the lifestyle of the user. Moreover in the UK, deferrable loads currently account for around 20% of the domestic electricity usage and this is likely to grow with the increased electrification of space and water heating [2, 9].⁵ Now, if agents are not coordinated in deferring these loads, they may induce higher peaks in the system than before (as discussed in Section 1). Thus, before proposing our solution to this, in the following subsections we build a model of a smart home that is sufficiently realistic to cap-

ment through market-based exchanges of energy rights to reduce peaks in electricity demand [15], and (iii) optimise the heating module of a building or a set of buildings through the use of spot pricing.

⁴See more details at <http://gridwise.pnl.gov/>.

⁵Mackay convincingly argues that heat pumps are much more efficient than (micro) combined heating and power (CHP) technology that is currently more popular.

ture the key effects that we address. In the next subsection, we elaborate on the types of devices we consider.

3.1 Deferrable Loads

Within the domestic energy domain, it is common to characterise the devices under four specific categories [5]: (i) wet (e.g., washing machine or dishwasher), (ii) cold (e.g., fridge or freezer), (iii) water heating (e.g., boiler or hot water dispenser) (iv) space heating (e.g., heat pumps or radiators). The devices that fall under these different categories behave very differently. For example, the wet types usually involve set periods of time at which they are switched on and off either by the user or by the device controller. The cold, water and space heating devices are much more dependent on the usage and the temperature of the home and will vary their behaviour depending on various external factors beyond the control of the user. Hence, we categorise these different loads in terms of those that have precise running times and running periods, which we call *shiftable static loads* (SSLs), and those that are more dependent on the temperature, which we call *thermal loads*.⁶ Given this, in the next sub-sections, we provide exemplar optimisation models for SSLs and a thermal load. We will consider a typical domestic profile that spans the usage over a day divided in half-hourly slots represented by a set of time slots $t \in T$ where $T = \{1, \dots, 48\}$. We also assume the smart meter receives a price signal (which may be fixed) at every time t .

3.2 Shiftable Static Loads

The set of SSLs is noted as $l \in L$ where each load is characterised by a set of parameters. Thus, $o_l \subseteq T$ is a set of time slots at which the device is set by the user to turn on. In effect, o_l represents the preferred time at which the user would like to switch on the device (e.g., switch on the washing machine while she is away). Any deviation from this preset time is likely to cause a loss of comfort to the user. Let $d_l \in \{-24, -23, \dots, 23, 24\}$ be the deferment of the device and note $r_l \in \mathbb{R}^+$ as the power rating of the device in kW. The duration v^l (in time slots) of each load is drawn from a uniform distribution, $U(v_{min}^l, v_{max}^l)$ where $v_{min}^l, v_{max}^l \in T$ will depend on the device used (e.g., a washing machine typically runs between 1 and 3 hours). To model the effect of shifting loads from their preset time, let $\Delta c_l \in \mathbb{R}^+$ be the *marginal comfort cost* associated with deferring the device to a different time from those initially set in o_l . Then, the overall comfort cost is $c_l = \Delta c_l |d_l|$. This assumes that the more the smart meter deviates from the time at which the user initially sets a device to run, the more discomfort it will cause to the user.⁷ The key variable in our optimisation model is effectively d_l since it determines the power consumption (which, in turn, determines the price paid by the user) and comfort cost of the user. Given a price p_t (in £/kWh), $\forall t \in T$, the objective is then to minimise the cost and comfort cost, as follows:

$$\begin{aligned} & \min_{d_l \forall l \in L} \sum_{l \in L} \sigma c_l + \sum_{t \in T} \sum_{l \in L} \gamma_t^l r_l p_t v^l 0 \\ & \text{such that } \gamma_t^l = \begin{cases} 1 & , t = t' + d_l \\ 0 & , \text{otherwise} \end{cases}, \quad (1) \\ & \text{where } t' \in o_l \end{aligned}$$

⁶It is possible to cast cold loads such as small fridges under SSLs if, for example, it is only required for the fridge to keep its contents well below room temperature (instead of a set temperature) by turning it on and off at regular intervals).

⁷Of course, more complex functions could be built by the user, but this is beyond the scope of this paper.

such that $\gamma_t^l \in \{0, 1\}$ determines when the devices is turned on or off (based on the shifting of the on times o_l by d_l), v_l is a number of time slots over which device l is on, and $\sigma \in [0, 1]$ is a scaling factor determining how much comfort is more important to the user relative to price. Using the above mathematical programming formulation, the optimisation problem of finding the optimal deferment of SSLs be therefore be directly solved using standard solvers.

3.3 Thermal Loads

We now turn to modelling the thermal properties of the home and the use of a heater controlled by an *intelligent* adaptive thermostat controlled by the smart meter.⁸ Standard adaptive thermostats warm up the house to be at the required temperature exactly when the user expects to be home. By turning on the heating before the user is present, the thermostat ensures the house is always warm when the user requires it to be. An intelligent adaptive thermostat, in turn, aims to optimise the heating (or cooling) profile (i.e., times at which the heater is turned on), in order to guarantee the required temperature is reached at the time the user is present *and* that the costs of doing so are minimised. To this end, we specify the properties of our heating optimisation model as follows. First, we detail the properties of the home and the heater. The heater we choose to model is a heat pump (either air-source or ground-source) which simply extracts heat from one place and moves it to another place [9]. Second, we provide a mixed-integer quadratic programming (MIQP) formulation of the problem which can be solved using off-the-shelf solvers.⁹

3.3.1 Home and Heater Properties

To model and simulate the thermal properties of the home and heater, we build upon the models proposed by [4] which use the ASHRAE model of comfort.¹⁰ Let $\phi \in \mathbb{R}^+$ be the thermal leakage rate of the house measured in W/K. The internal temperature of the home at time t is $\tau_{in}^t \in \mathbb{R}^+$ and the external temperature (in K) is denoted as $\tau_{ext}^t \in \mathbb{R}^+$. The user will set the temperature $\tau_{opt} \in \mathbb{R}^+$ at the level she feels comfortable (typically 22.5 degrees Celsius). The heat capacity and the total mass of air in the home are denoted as $a_{hc} \in \mathbb{R}^+$ (in J/kg/K) and $a_{mass} \in \mathbb{R}^+$ (in kg) (which is computed using the air density and house volume) respectively. The heater is described by r_h and o_h where $r_h \in \mathbb{R}^+$ is the power rating (in kW) of the heater and $o_h \subseteq T$ is a set of time slots at which the heat pump is switched on. Given this, we also define the variable $h_{on}^t \in \{0, 1\}$ for every $t \in T$ where $h_{on}^t = 1$ if $t \in o_h$ and 0 otherwise. This description of the heater is similar to the deferrable loads above except that the deferment of the heater is much more complex as we will see next. For now, we can compute the total heat input in the system as:

$$\eta_i^t = h_{on}^{t-1} r_h - \phi(\tau_{in}^{t-1} - \tau_{ext}^{t-1}) \quad (2)$$

Moreover, we can define the relationship between the inside temperature at time t and the amount of heat injected in

⁸The model also applies to cooling effects simply by setting the parameters differently to model heat extraction rather than heat injection as we do here.

⁹We use IBM ILOG CPLEX 12.2 in our experiments.

¹⁰American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Standard 55-2010: Thermal environmental conditions for human occupancy.

the house at time $t - 1$ as:

$$\tau_{in}^t = \tau_{in}^{t-1} + \frac{k\eta_i^t}{a_{hc}a_{mass}} \quad (3)$$

where $k = \frac{24 \times 3600}{|T|}$ is the time (in seconds) during which the heater is on.¹¹

Now, we expect users to set the times at which they will be at home and will want the temperature adjusted for their comfort. To this end, we define comfort on times as a $C_{on}^t \in \{0, 1\}$ at every time slot $t \in T$ such that the $C_{on}^t = 1$ if the user needs comfort and $C_{on}^t = 0$ otherwise. During every time slot, we model the *instantaneous* comfort cost (i.e., ignoring changes from the previous time slot) of the user due to heating as $\Delta c_h^t \in \mathbb{R}^+$ such that:

$$\Delta c_h^t = \begin{cases} C_{on}^t \omega_1 (\tau_{in}^t - \tau_{opt})^2, & \tau_{in}^t \geq \tau_{opt} \\ C_{on}^t \omega_2 (\tau_{in}^t - \tau_{opt})^2, & \tau_{in}^t < \tau_{opt} \end{cases} \quad (4)$$

where $\omega_1, \omega_2 \in [0, 1]$ are constants that scale the effect of the temperature difference depending on whether it is more (or less) comfortable when the temperature is higher or lower than t_{opt} (in degrees celcius). Typically, in colder periods, $\omega_1 < \omega_2$ is preferred. Then, the total comfort cost at time t is a combination of the present instantaneous comfort cost and at $t - 1$ given by:

$$c_h^t = \Delta c_h^t + \gamma \Delta c_h^{t-1} \quad (5)$$

where $\gamma \in [0, 1]$ scales the effect of the previous time slot on the current one (and captures the psychological persistence of discomfort).

3.3.2 MIQP Formulation

Given the times when $C_{on}^t = 1$ and the comfort cost c_h^t , the goal of the agent is to optimise the heating so as to minimise the price paid by the user. Thus, assuming the price of electricity at every time t is p_t , and specifying $h_{on}^t \in \{0, 1\}, \forall t \in T$ as the decision variables, the objective function of the program is:

$$\min_{h_{on}^t, \forall t \in T} \sum_{t \in T} c_h^t + \kappa (v^h h_{on}^t r_h p_t) \quad (6)$$

subject to (2),(3), and (4), where $\kappa \in [0, 1]$ balances the cost of heating against the comfort and $v^h = 1800s$ is the duration of one time slot in seconds (assuming each time the heater is turned on for a half-hour). The quadratic nature of the objective function arises as a result of the computation of the comfort cost in (4). We set κ to a very small value to ensure that this part of the objective function only ensures that the user mainly optimises her comfort. Now, in some cases, a budget can be set by the user according to how much she can afford to spend on a daily basis for heating purposes. Thus, in order to ensure that the cost is never higher than a set value $p' \in \mathbb{R}^+$, we also add the following constraint to the model:

$$\sum_{t \in T} v^h h_{on}^t r_h p_t \leq p' \quad (7)$$

Since agents are only intent on maximising their gains (in comfort and cost savings), their aggregated uncoordinated behaviour may result in poor system performance unless demand-side management mechanisms are put in place.

¹¹In our experiments, we refine the resolution to smaller time slots (e.g., of 5 or 10 minutes) to get a better measure of the temperature inside the room.

Hence, given our model of the smart home, in the next section, we first describe the optimal and centralised solution that maximises social welfare (and describes what optimal means in this context) that determines the optimal coordinated behaviour of the agents. By so doing, we can use this as a benchmark for any control mechanism that can be developed in this domain (which we use to evaluate our mechanism in Section 6).

4. MAXIMISING SOCIAL WELFARE

As a result of the individual agents' deferment of loads, the consumption of energy at different times of the day may significantly increase or decrease, resulting in price spikes in the system. Before going on to design a control mechanism to reduce this effect, however, we need to determine the optimal performance of the population of agents as a whole to evaluate our control mechanism (see Section 6).

As was shown in [13], when agents attempt to defer consumption (in their case, they used storage), a game theoretic analysis of the problem predicts that the agents should converge to the competitive equilibrium. This means that they converge to an equilibrium (flat) price for electricity if there is enough storage, at which point social welfare is maximised.

In the context of deferrable loads, such an analysis is significantly more complex as shiftable static loads, comfort costs, and thermal loads are not homogeneous across the population of agents. Instead, we can compute the point at which all agents behave optimally by aligning their individual objective functions with the global optimum as follows. First, we assume that each agent $i \in I$, where I is the set of agents, has a non-shiftable static load l_{fixed} such as lighting, cooking and other types of essential loads which have a power rating $fixed_{t,i} \in \mathbb{R}^+$ in kW at every time point $t \in T$. Second, we assume that the cost of electricity for the population is given by a quadratic cost function (the function is assumed to be quadratic to mimic an increasing marginal cost for electricity supply – we use such a model in our evaluation) $s_t : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ for every time slot t such as $s_t(q) = \theta q^2$ where $\theta > 0$. This means that the greater the demand from the agents, the higher will be the unit price of electricity (because of the monotonically increasing function). Then, we aggregate the functions in equations (1) and (6) subject to the individual constraints as stated in Equations (2), (3) and (5) and the conditions stated in Sections 3.2 and 3.3.1 as follows:

$$\min_{h_{on}^t, \forall t \in T, d_i^t, \forall i \in I} \sum_{i \in I} \left(\begin{aligned} & \sum_{t \in T} c_{h,i}^t + \sum_{l \in L_i} \sigma c_{l,i} \\ & + \sum_{t \in T} s_t \left(\sum_{l \in L_i} \gamma_{t,i}^l r_{l,i} v_i^l \right) \\ & + \kappa v_i^h h_{on,i}^t r_{h,i} + fixed_{t,i} \end{aligned} \right) \quad (8)$$

The main difference between the above objective and the individual agents' objective is that the cost of electricity in the above case is based on the induced cost as a result of the agents' aggregated demand (both from their deferrable loads and their static demand) using the supply function $s(\cdot)$ rather than the cost predicted by the agents p_t (which is assumed to be a fixed unit cost).

Note that the optimal algorithm above is unlikely to scale very well in the number of loads per agent and deferability of these loads given the complexity of the optimisation problem and given non-linearity of the objective function. Specifically, experiments (where Equation (4) is linearised) using a standard MIQP solver (e.g., IBM ILOG CPLEX

12.2) could be run for up to 75 agents with up to three deferrable loads each (on a 64-bit machine with 12GB of RAM). Beyond this size, the solver cannot model the problem in memory. This would be fine for scenarios where a central controller has complete control over the few deferrable loads in a building for example, it will clearly not scale to hundreds of houses (where users might not want the central controller to override their preferences to use electricity) or for deferrable loads (and certainly not to all 26M UK homes).

Hence, in the next section, we discuss some of main issues in DSM and introduce a novel approach to tackling this problem in a completely decentralised fashion.

5. DECENTRALISED DEMAND-SIDE MANAGEMENT

Most DSM approaches involve a central controller that advises a pool of consumers to reduce their current demand (during peak demands), often by reducing or deferring loads subject to economic incentives. This approach has repeatedly been shown to be effective for relatively small pool sizes of industrial and commercial consumers [6]. Indeed, it constitutes the business model of a number of energy service companies (ESCOs) that exclusively deal with Demand Side Management (e.g., EnergyConnect Inc. and EnerNOC). While it remains feasible to signal a small number of consumers and expect an immediate response, DSM at a regional or national level (with 26M of households in the UK) is more complex. Given this, in the next subsections, we discuss the main issues associated with applying DSM on a large scale (dealing with thousands and millions of homes). First, we justify the need for a price signal to incentivise the agents to defer their demand. Second, we show that, even if an accurate price signal is provided according to the criteria we propose, the adaptive and autonomous behaviour of the agents in the system is the key component that can enable significant performance benefits in the Smart Grid. Thus, we propose a novel model of decentralised demand side management (DDSM).

5.1 The Pricing Mechanism

The pricing mechanism that we propose involves signalling the costs of generating electricity to the consumers. Traditionally, consumers are offered a fixed price for electricity by their supplier. Exceptions to this include time-of-use (TOU) pricing (e.g., Economy 7 heating or eco:2020 in the UK) and real-time pricing (RTP) schemes. While fixed pricing mechanisms completely hide the real-time costs of electricity (which varies according to the type of generators used and availability of intermittent renewable energy), TOU pricing simply *bias*es the real price of electricity in order to incentivise users (who typically aim to maximise their savings) to shift their loads to off-peak periods (i.e., when aggregate demand is lower). In contrast, RTP involves providing a price signal for the next 30 minutes time slot at the current time. While RTP is still being evaluated in a number of trials, it has been shown to be better than TOU pricing in allowing users to dynamically adjust their demand to avoid peaks when electricity is more expensive [6, 7].

To understand the effect of the different pricing schemes on demand, consider Figure 1 where we show how 500 smart homes¹² (as described earlier) would react to the three dif-

¹²We implemented each smart home according to the settings

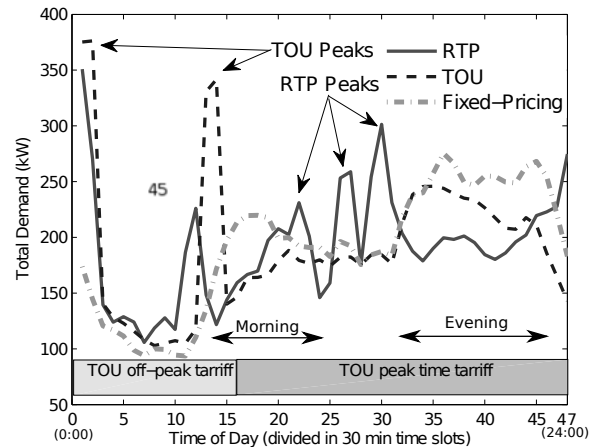


Figure 1: Demand optimisation given the Domestic Economy 7 tariff and the 30-min-tariff pricing models.

ferent types of pricing mechanisms (where real time prices are provided for 30-minute time slots). As expected, a fixed price induces a relatively high demand with large plateaus in the morning and evening (when users are at home and typically use their loads) since users are indifferent about the real-time costs of generating electricity. More importantly, we can observe that when all the consumers are on a two-tariff mechanism, they optimise their demand at the same time (to take opportunity of off-peak prices), and thus the ensuing demand on the grid now peaks during off-peak times. Similarly, it can be seen that the RTP mechanism also shifts demand to different times of the day and induces other peaks at these times.

These results show that a more accurate signal (i.e., representing real costs) allows consumers to better optimise their demand by reacting more often to a more accurate 30-min-tariff pricing model. However, it is also clear that the demand *cannot be flattened* by applying *only* a RTP mechanism, while completely ignoring the behaviour of the agents. This is because if the agents are signalled a low price for the next 30-minute period, they will all switch on their devices, which then results in a peak in demand at the next time period. When such a mechanism is rolled out on a large scale, such *reactive* behaviours can cause significant peaks. To remedy this, in the next sub-section, we propose a novel adaptive behaviour for the agents, which builds upon the RTP mechanism, to allow agents, *without any centralised control*, to coordinate in our DDSM model.

5.2 The Adaptive Mechanism

Our adaptive mechanism for DDSM is composed of algorithms that determine how to defer the two different types of deferrable loads present in the home (i.e., shiftable static loads and thermal loads) and how to optimise heating. First, we adopt the Widrow-Hoff learning mechanism [13] to gradually adapt how agents defer their deferrable loads based on predicted market prices for the *next day*¹³ (as opposed to the next 30 minutes as in [6, 7]). Specifically, an agent i gradu-

given in the evaluation section.
¹³We use a weighted moving average to predict day-ahead market prices, but more sophisticated approaches incorporating domain-specific knowledge could easily be used.

ally adapts its deferment parameter d_i^i towards the optimal $d_i^{i,*}$ (as computed by the objective function in Equation (1)) as follows:

$$d_i^i(t+1) = d_i^i(t) + \beta^i(d_i^{i,*} - d_i^i(t))$$

where $\beta^i \in (0, 1]$ defines the learning rate (i.e. how fast the agent reacts to changing conditions).

Second, each agent’s behaviour is further modelled by how often it readjusts its heating profile. Specifically, it reoptimises its thermal load profile on any particular day, with a probability of $\alpha \in (0, 1)$. This means, it reruns the optimisation detailed in Section 3.3.2 with a probability of α according to the predicted prices p_t for the next day. The expected number of agents that will optimise on any particular day is thus given by $N\alpha$, where N is the number of agents with DSM capability. By so doing, we introduce inertia into the optimisation by ensuring that all DSM-capable agents do not reoptimise at the same time and, in the next section, we empirically demonstrate that a stable and desirable outcome is reached whereby grid demand flattens when α and β are sufficiently small. Moreover, we also show that the agents’ best strategy (profit maximising) is to adopt an adaptive behaviour. Thus, even though consumers are not centrally coordinated for DSM, our mechanism ensures an effective *emergent* coordination among consumers.

In general, our DDSM mechanism can be particularly attractive when dealing with millions of consumers since it does not require, apart from the price signal (which could be directly sent to the smart meter), any other form of direct communication between the generators and the consumers. Moreover, the adaptive mechanism can also be modelled for different types of consumers as discussed above. Thus, when large populations of such consumers are simulated within a smart grid, it is possible, with reasonable accuracy, to predict the behaviour of the system. To this end, in the next section, we provide simulations of populations of smart homes (extrapolated across the grid¹⁴) and evaluate the performance of our DDSM both in terms of the consumers’ benefits and the grid performance. In so doing, we aim to show how our DDSM mechanism generally aligns the incentives of the users with the system-wide objectives, leading to a more efficient and stable Smart Grid.

6. EMPIRICAL EVALUATION

In our experiments we consider a population of 5000 agents¹⁵ and a RTP pricing based on the macro-model of the UK electricity market (as outlined by Vytelingum *et al.* [13]), with a quadratically increasing marginal cost. We used, for the shiftable static loads and thermal loads, data compiled by the Department of Trade and Industry in the UK from 2008.¹⁶ This dataset shows that, while shiftable static loads

¹⁴We use a similar methodology to [13] based on real UK Grid demand and market prices and average UK domestic load profiles for a typical weekday in winter.

¹⁵The time taken to run our simulations grows linearly with the number of agents and we have confirmed similar results with populations of up to 10,000 agents. However, when we compare to the optimal solution, we are limited to 75 agents, due to the computational complexity of the solution. Hence, to approximate the optimal solution for 5000 agents, we ran the global optimisation in Equation (8) 100 times for numbers of agents from 50 to 75 and estimated the optimal solution for 5000 agents from the trend identified.

¹⁶<http://www.berr.gov.uk/files/file11250.pdf> on 20/10/2010.

(e.g., wet loads such as washing machines and dryers) are owned by nearly the whole population, the penetration of electrical heating is currently only around 7%. Moreover, shiftable static loads take up on average 20% of the total energy usage of a house while the rest of the consumption is due to entertainment, lighting, and cooking purposes. This setup is likely to change in the future because of a growing need to use more efficient heating devices such as heat pumps such that the load profile is likely to change significantly.

Given this, in this section, we also aim to show the effects of increased deferment of electrified thermal load, on the smart grid. To this end, we create instances of the smart home with parameters set according to the dataset above. We model two wet shiftable static loads (dryer and washing machine) with power rating of 2kW each as well as a thermal load (with $\tau_{opt} = 22.5$, $a_c = 1000$, $a_{mass} \in [1000, 2000]$, and $r_h = 2kW$). In order to generate the specified ON times for the shiftable static loads, we use the average load profile as a probability density function (i.e., describing the probability of the device being set to be switched on by the user at each half-hour period) and use a Poisson distribution $Pois(\lambda)$ to generate the number of times the devices are switched on in a day (with $\lambda = 2$). Moreover, we generate τ_{ext} as the average 5 minutes temperature readings taken over a period of 5 days from the roof-top sensor at one of our university buildings in January 2010 (during the cold season). Wherever needed, we adjust the number of agents having electrified heating (i.e., starting with the current 7%). Our experiments are repeated 100 times and the results averaged and error bars plotted to represent the 95% confidence intervals.


6.1 Performance of the Adaptive Mechanism

First, we evaluate the efficiency of our mechanism. We do so by comparing the efficiency of the grid when using our DDSM mechanism as opposed to a centralised system with complete and perfect information against an optimal behaviour. Given our optimal solution (described in Section 4) with an optimal load factor¹⁷ (LF) of 0.76, we can observe that a DDSM-based grid behaviour converges to the optimal behaviour (up from LF=0.6) when α is reasonably small as shown in Figure 2(a). While higher values may give faster convergence to the optimal, they do not allow the system to settle at a more efficient equilibrium (i.e. closer to the optimal). Furthermore, when α is too large, there is no convergence as too many agents are reoptimising at the same time such that the peaks are simply moved rather than flattened. When $\alpha = 1$ and every user optimises at the same time, the system breaks down (with an LF averaging 0.55). Thus, as we empirically showed, the system converges to the optimal behaviour without any centralised coordination when α is reasonably small. Specifically, for the rest of our experiments, we set $\alpha = 0.05$.


6.2 Emergent Behaviour of the Smart Grid

Given our DDSM approach with its adaptive mechanism, we first empirically demonstrate in Figure 2(b) that as the load factor of domestic consumers converges to the optimal load factor at 0.76 (meaning fewer peaks in the system), the

¹⁷The load factor is the ratio of average power to peak power and is ideally 1. A low load factor suggests large peaks in demand. Here we computed the load factor based on our estimation of the optimal solution as discussed in the previous section. We validated our results on small numbers of agents (50-75).



(a)



(b)

Figure 2: Evaluating DDSM in the Smart Grid w.r.t. (a) learning rates, (b) Load Factor & CO₂ emissions.

percentage of carbon reduction also increases up to 6% (i.e., electricity is produced from less polluting sources) as fewer carbon intensive peaking plants are used.

Next, we analyse the system by considering the load duration curve (LDC) of the grid. The LDC¹⁸ is used to illustrate the relationship between generating capacity requirements and capacity utilization, normalised to the current domestic load demand in UK. Figure 3 shows the LDC of the system (based on today's 7% electrification of heating) and, specifically, we can observe that particularly within the peak-load region (which is from 100% to 80%), the curve flattens when optimised, which implies a flattening of demand peaks. Furthermore, at 100% peak demand, we observe that the domestic load decreases by up to 17%, compared to the unoptimised case, which implies that the grid requires 17% less capacity to cope with domestic demand. Considering that current domestic peak demand is in excess of 15 GW, a 17% reduction in capacity requirement is significant and will become increasingly vital given the high rate of increase of domestic peak demand.

Next, we analyse the population dynamics to evaluate whether users will be incentivised to automate their smart meters (i.e., to give control to an agent). Specifically, we use evolutionary game theory (EGT) techniques which allow us to determine whether different proportions of the population will adopt smart meters (or not) depending on how much they can save by doing so. First, we formulate the problem as a complex non-deterministic game where agents have a mixed strategy $x_r \in (0, 1)$, i.e. a probability that they have DDSM capability and are only motivated by financial gains, where $r \in S = \{DDSM, -DDSM\}$. By analysing how x_r changes as the payoffs of agents with and without DDSM change for different x_r (assuming agents are more likely to




Figure 3: Load Duration Curve of Domestic Consumers with today's 7% electrification.

adopt the better strategy – whether or not to adopt smart metering), we aim to analyse how the proportion of the population adopting smart meters and DDSM evolves using the following equations [14]:

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

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

where $u(x, x)$ is the expected payoff of any agent (whether using DDSM or not), $u(e^r, x)$ is the expected payoff of an agent with DDSM given a proportion of the population with DDSM of x_r and, finally, x_{nash} is the Nash Equilibrium of the system (i.e., x_r where there are no incentive for an agent to deviate from). Figure 4 shows the payoffs of agents with and without DDSM. Based on our EGT analysis, we deduce that there is always an incentive for an agent to adopt DDSM given the higher payoffs for any x_r , such that \dot{x}_r is always positive. The population dynamics eventually converges to the Nash Equilibrium, $x_{nash} = 1$. This implies that all agents eventually adopt smart meters and DDSM.

¹⁸The load duration curve is computed as follows. Let ℓ be a vector of load values of the system at different time slots and ℓ^* be the ordered version (from high to low) of ℓ . The load duration curve is then given as $(f(x) \in X = \{\ell_1^*, \dots, \ell_M^*\} : x \in T^+ = \{T_1^+, \dots, T_M^+\})$, where $T_i^+ = \sum_{j=i}^M t_j$ and t_j is the time slot of the j^{th} ordered load value. The left-hand side represents the peak loads of the system. The LDC is a useful way of breaking down the load factor value to describe how long peaks last for. The domestic capacity (peak load) is given at 100% load duration.




Figure 4: Percentage savings of agents with and without DDSM for different proportions of the population adopting smart meters and DDSM.




Figure 5: Load Duration Curve of Domestic Consumers with 25% electrification.

6.3 The Effect of Electrical Heating

Finally, we analyse future scenarios where the proportion of smart homes with thermal loads might increase. Specifically, Figure 5 shows grid performance for a potential future of 25% electrification. Thus, if the number of smart homes with thermal loads were to increase from 7% to 25%, the demands would be much higher (by up to 31% – based on initial loads from Figure 3), with higher peaks. Using DDSM, we demonstrate that these peaks would again be flattened significantly (shown by the significant drop in the 100% to 80%-peak region) using our novel model of a DDSM, with a 22% decrease in peak domestic capacity.

7. CONCLUSIONS

In this paper, we have provided a model of a smart home and presented our DDSM; a novel paradigm for mechanisms to manage demand on a large scale in the smart grid. Through our simulations involving 5000 homes and using average (winter) load profiles for 26M homes in the UK, we have shown that our DDSM mechanism can improve grid per-

formance by reducing peaks in demand by up to 17% and carbon emissions by up to 6%. We have also predicted, using evolutionary game theoretic techniques, that consumers will have significant economic incentives to adopt agent-based smart meters and will eventually all do so. Finally, we have shown that DDSM also reduces peaks as demand grows as a result of increased electrified heating. Future work will look at integrating price and weather predictions for more effective DDSM mechanisms and extending these mechanisms for commercial and industrial consumers in the smart grid. We also aim to evaluate our mechanism in scenarios where not all buildings are equipped with agents that can optimise their behaviour and remain insensitive to real-time prices as it is at present.

Acknowledgments

The work presented in this paper is funded through the IDEAS project (<http://www.ideasproject.info>) by a company active within the energy domain.

8. REFERENCES

- [1] Department of Energy and Climate Change. Smarter grids: The opportunity. HM Government, 2009.
- [2] Department of Energy and Climate Change. The UK low carbon transition plan: National strategy for climate and energy. Technical report, UK Dept. of E & CC, 2009.
- [3] R. Galvin and K. Yeager. *Perfect Power: How the MicroGrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill Professional, 2008.
- [4] Y. Guo, R. Li, G. Poulton, and A. Zeman. A Simulator for Self-Adaptive Energy Demand Management. In *Proc. of the 2nd IEEE Conf. on Self-Adaptive and Self-Organizing Systems*, pages 64–73, 2008.
- [5] V. Hamidi, F. Li, and F. Robinson. Demand response in the uk domestic sector. *Electric Power Systems Research*, 79(12):1722 – 1726, 2009.
- [6] D. Hammerstrom et al. Pacific Northwest GridWise Testbed Demonstration Projects; Part II. Grid Friendly Appliance Project. Technical report, PNNL-17079, Pacific Northwest National Laboratory, 2007.
- [7] D. Hammerstrom et al. Pacific Northwest GridWise Testbed Demonstration Projects; Part I. Olympic Peninsula Project. Technical report, PNNL-17167, Pacific Northwest National Laboratory, 2008.
- [8] D. G. Infield, J. Short, C. Home, and L. L. Freris. Potential for domestic dynamic demand-side management in the uk. In *Proc. of the IEEE Power Engineering Society General Meeting*, pages 1–6, 2007.
- [9] D. MacKay. *Sustainable energy without the hot air*. UIT, Cambridge, 2009.
- [10] F. Scheppe, B. Daryanian, and R. Tabors. Algorithms for a spot price responding residential load controller. *IEEE Power Engineering Review*, 9(5):49–50, 1989.
- [11] F. Scheppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox. Homeostatic utility control. *IEEE Transactions on Power Apparatus and Systems*, 99(3):1151 – 1163, 1980.
- [12] US Department of Energy. Grid 2030: A National Vision For Electricity’s Second 100 Years, 2003.
- [13] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. In *Proc. of the 8th Conf. Autonomous Agents And MultiAgent Systems, AAMAS 2010*, pages 39–46, 2010.
- [14] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.
- [15] R. D. Williams and F. C. Scheppe. Peak-power-demand limitation through independent consumer coordination. *IEEE Transactions on Systems, Man and Cybernetics*, 16(4):560 – 569, 1986.