

Accommodating renewable generation through an aggregator-focused method for inducing demand side response from electricity consumers

P.J. Boait, B. M. Ardestani, J.R. Snape

Institute of Energy and Sustainable Development, De Montfort University, Queens Building, The Gateway, Leicester LE1 9BH, UK.

e-mail: p.boait@dmu.ac.uk

Abstract

The ability to influence electricity demand from domestic and small business consumers so that it can be matched to intermittent renewable generation and distribution network constraints is a key capability of a smart grid. This involves signalling to consumers to indicate when electricity use is desirable or undesirable. However simply signalling a time dependent price does not always achieve the required demand response and can result in unstable system behaviour. We propose a demand response scheme in which an aggregator mediates between the consumer and the market and provides a signal to a “smart home” control unit that manages the consumer’s appliances using a novel method for reconciliation of the consumer’s needs and preferences with the incentives supplied by the signal. This method involves random allocation of demand within timeslots acceptable to the consumer with a bias depending on the signal provided. By simulating a population of domestic consumers using heat pumps and electric vehicles with properties consistent with UK national statistics, we show the method allows total demand to be predicted and shaped in a way that can simultaneously match renewable generation and satisfy network constraints, leading to benefits from reduced use of peaking plant and avoided network reinforcement.

1. Introduction

There is wide acceptance that a key element of the smart grid will be the ability of the network to provide a price or cost based signal to domestic and small commercial or industrial electricity consumers that will elicit a useful response from them. This response will comprise variations to demand and/or micro scale generation resulting from an elastic response to price or

displacement of demand or generation, arising through deferment of appliance use or exploitation of energy storage available to the consumer. The potential benefits sought from this feedback loop are twofold.

Firstly, an increased ability on the consumption side to track variations in output from low carbon renewable generators and thereby minimise the need for high cost/high carbon standby generation. The UK's Renewable Energy Roadmap [1] predicts that installed wind generation capacity in the range 21-31GW will provide between 13 and 20% of supplied electricity in 2020. Projections to 2050 [2] envisage this rising to 35%. There will be a corresponding need within the UK electricity system to manage the variability of wind generation output which Gross *et al* [3] have shown implies additional despatchable reserve capacity of up to 10% of the wind capacity, potentially adding £3-£5/MWh to the cost of wind generation. Since average UK domestic electricity consumption in 2010 was 13.6GW [4], the ability to vary this at will would make a valuable contribution to the additional reserve requirement. This paper will show how, given a suitable control system integrated with the national smart metering network in about 70% of electrically heated homes, domestic demand response could provide the equivalent of 0.7-2.8 GW of despatchable capacity in 2020 representing up to 90% of the additional reserve requirement needed to support planned wind generation capacity at that date.

Secondly, reduction in power flow peaks in congested elements of the electricity distribution infrastructure is necessary so that increases in demand from electrification of space heating and transport can be accommodated without massive investment. The 2050 scenario in [2] indicates annual electricity consumption of 883 TWh, more than twice the current level. The potential value of a capability to mitigate local demand peaks is quantified by Pudjianto *et al.* in [5] at up to £30Bn through deferred or avoided network reinforcement costs.

A variety of schemes for construction and operation of this feedback signal have been tried in practice or by simulation and no generally accepted paradigm has emerged. This paper begins

with a review of the experience to date and highlights a critical weakness in the system stability of approaches based on direct response to a time-varying price of electricity. In section 3 we describe a solution to this problem comprising two elements; appropriate mediation between the wholesale electricity market and the consumer by an aggregator who holds the retail electricity supply relationship with the consumer, and a “smart home” control unit which manages the consumer’s appliances using a novel algorithm for reconciliation of the consumer’s needs and desires with the incentives supplied by the signal. Section 4 sets out the assumptions employed in a simulation to test this proposal, based on UK national statistics for appliance and vehicle ownership and use. Results are then presented showing the ability of this approach to provide both exploitation of renewable generation availability and despatchable reserve capacity. Sections 5 and 6 discuss the implications of these results and draw conclusions.

2. Background

Domestic demand response schemes typically consist of some combination of electricity tariffs which incentivise response, a signal to the consumer indicating when the response should take place, and agency in the form of manual or automatic control of electricity use that executes the response. An illustrative example is the UK “Economy 7” scheme which offers a discounted tariff for 7 hours overnight and provides a signal embedded in a 198 kHz radio broadcast that indicates the start of the 7 hour interval and causes electric heating appliances with thermal storage to switch on. The population of consumers with this facility (about 1 million) is divided into groups receiving a different start time to ensure this demand is ramped up and distributed in time to satisfy balancing requirements from the system operator. In addition these consumers often set their own time clocks such that water heating or wet appliance operation is performed during the discounted period. The system therefore combines automatic and human agency.

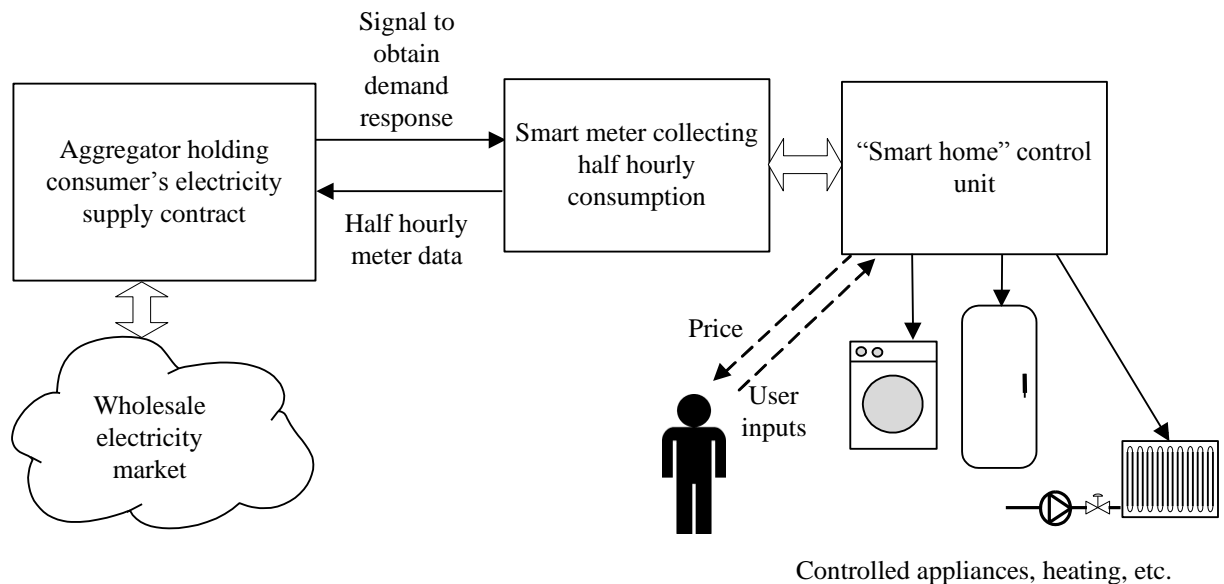


Figure 1. System for inducing consumer demand response

Research on the effectiveness of smart meters and dynamic pricing of electricity indicates that:

- the willingness of consumers to respond manually to a display of their electricity cost without dynamic pricing is limited and tends to decay with time [6];
- time-of-use pricing alone can consistently reduce peak consumption – a range of US trials evaluated by Faruqi and Sergici [7] show between 2%-6%;
- introduction of some “enabling technology” to perform the response results in much higher peak reductions (21%-32% in [7]).

These results, particularly the third point, imply that the main burden of demand response should be undertaken by some form of automatic control system which responds flexibly to a price-based signal allowing a wide range of appliances to be brought under control at the consumer’s discretion. This conclusion has motivated numerous recent studies of possible schemes by which domestic appliances and electric vehicles might respond automatically to a price signal [8],[9],[10]. Figure 1 illustrates the main features and information flows of this type of system. A common feature of all these investigations is that the control unit executing the

response has an objective function to minimise cost for the consumer while satisfying the various constraints they may have set such as the latest time that the dishwasher must run. This allows the modelling to show that the consumer would save money and thereby have an incentive to participate in the scheme. However, cost minimisation as an objective function has a strong tendency to cause demand peaks to arise at times when the cost signalled is at a minimum. This can cause the peak-to-average ratio (PAR) of the resulting aggregate demand to exceed that of the baseline (i.e. demand with no feedback signal) and hence a key objective of demand response, to reduce PAR, is not delivered. Given the potential significance of this problem in any practical implementation it is worth reviewing these studies in more detail to see how the authors recognise it and the mitigations proposed.

Mohensian-Rad & Leon-Garcia [11, figures 8 and 11] show explicitly that for a household with a baseline mean daily PAR of 4.49 the effect of introducing half-hourly pricing and cost optimised control of heating and appliances is to increase mean PAR to about 9. The mitigation they introduce is inclining block rate (IBR) tariffs which imply that if household electricity consumption in a half hour interval exceeds a certain level then subsequent units are charged at a higher rate. In other words, the tariff then has both time of day and volume dependencies. When the home control unit optimises against this combined tariff PAR falls to 2.75. Ramchurn *et al.* [12, figure 1] show the peaking effect of time dependent pricing and cost optimising control and conclude that “demand cannot be flattened by applying only a real time pricing mechanism”. Their corrective is to ensure that home control “agents” do not optimise their response daily, but with a probability α in the range 0.05 - 0.2. This is effective in reducing the PAR of aggregate demand but such a scheme would not be able to respond to short term variations in supply in a timescale of less than $1/\alpha$ days, for example in wind generator output. Rastegar *et al.* [13] observe the peaking effect so they constrain cost optimisation with limits on the displacement of demand in time and on total consumption within a timeslot. Roscoe and Ault [9] identify the peaking problem very clearly (their figures 9 and 10) and because their modelling includes a

market model they show oscillation in market prices resulting from the gain in the price-demand feedback loop. They suggest (but do not test) a variety of mitigations including intelligent demand forecasting, issuing different price profiles to different customers, and mandatory randomisation of response by the home controller. Randomisation was also proposed by Boait *et al.* [8] and is the kernel of the solution set out in this paper.

In essence, the existence of a direct feedback loop between the wholesale electricity market and a large number of automated cost optimising consumers creates a complex system with the potential to be unstable under a range of scenarios. Instead of patchwork mitigations to this fundamental issue, the first premise of the solution proposed here is that cost minimisation is not performed in the home by a control unit operating on behalf of the consumer, but by the organisational entity holding the electricity supply contract with the consumer. We henceforth refer to this entity as the aggregator, since this term is often used e.g. by Giordano and Fulli [14] to describe a commercial organisation that has the ability to control or influence the electricity demand of a number of consumers (currently often industrial) and uses it to participate in an electricity market. However in the role envisaged here, the aggregator will also supply electricity to consumers and procure microgeneration exports from them, so performing the functions of a retail supplier. The aggregator sends a 48-valued day-ahead signal to the home control unit of its customers and the unit mobilises a predictable response based on the signal from its dependent appliances according to algorithms described in the next section. The signal is not necessarily a wholesale price, or any other price (although it can be interpreted as a price to the consumer as will be explained). It is determined in accordance with the economic interest of the aggregator; the tariff offered to the consumer is not constrained.

Where, as in the UK, there is a competitive market for retail electricity supply this should ensure that the consumer shares in any efficiency gains from demand response achieved by the aggregator. Because the scheme proposed here allows demand to be shaped to match available

supply and reduces risks of exposure to the balancing market, it is of particular potential value to, and should aid the viability of, local consumer-owned energy co-operatives which are of interest to policymakers for their ability to mobilise change towards CO₂ reduction targets and also address fuel poverty [15]. There are other advantages in ownership of the demand response capability by aggregators:

- a. Because benefits accrue in the first instance to aggregators, there is a better chance of a viable business case being constructed that will support roll-out of the capability as an add-on to the smart meter programme. Fragmentation of the demand response value chain has been identified as a barrier to adoption [16].
- b. Aggregators will have an incentive to accept the system integration cost and risk involved in configuring and maintaining a device that will interface between the smart meters and heating controls, appliances, and electric vehicles. Their role will become similar to that of internet service providers who support the setup of home routers and WiFi. Without this commitment, skill shortages are liable to impede take-up [17].

To incentivise consumers to participate in this scheme, the aggregator can use the smart meter data from each consumer to evaluate the correlation between their demand and the signal provided. A variable discount to a flat tariff could then be offered that depended on the level of demand response actually achieved, reflecting the number and demand characteristics of the appliances each consumer has engaged with the home control unit and any manual control actions they have taken that were synchronised with the signal. Other incentive structures are possible such as a direct relationship between the signal and price. A benefit of this concept is that it allows flexibility in retail electricity contracts implementing it.

3. Proposed scheme for inducing stable and predictable demand response

3.1 The home control unit

The function of the proposed domestic control unit is to distribute adjustments to the electrical demand of its controlled appliances over the 48 timeslots in each day in inverse proportion to the 48-valued signal S provided by the aggregator at midnight for the coming day. The signal is framed by the aggregator and interpreted by the control unit such that high values deter electricity use, and low values attract use, so that the signal can represent a cost. The response to the signal in each timeslot is arranged, as far as possible within the constraints on individual appliances, to be proportional to the deviation from the mean. Since to a first approximation any demand displaced from a given timeslot must be recovered in another, the values of S in normalised form must be constrained as:

$$-1 \leq S_i \leq 1 \quad \text{for } i = 1:48 \quad (1)$$

$$\sum_{i=1}^{48} S_i = 0$$

This proportionality is achieved by displacing demand to alternative timeslots within a time window acceptable to the consumer with a probability proportional to the amplitude of the signal in each timeslot of the window. Because the consumer's demand must be met at some time during the time window, the probabilities of demand occurring in individual timeslots within the window must add up to 1. This means that for any S with more than one non-zero timeslot within the window some redistribution of demand will take place giving this method an inherent tendency to reduce peaks in demand. Taking the example of a wet appliance that the consumer has specified must run sometime in the next n timeslots with a duration of t timeslots, the

controller will create an “attractiveness” vector¹ A_w from the signal S with values given by $A_{wi} = (S_{max} - S_i)$ where S_{max} is the largest value of S . The controller then uses a random number generator to select one of the $n-t+1$ possible timeslots in which the appliance can start, with a probability P_{wi} given by:

$$P_{wi} = \frac{\sum_{j=i}^{i+t-1} A_{wj}}{\sum_{q=1}^{n-t+1} \left(\sum_{r=q}^{q+t-1} A_{wr} \right)} \quad (2)$$

The effect is that over a population of consumers choosing a similar time window demand will be allocated in proportion to the possible alternative averages of A_w over t timeslots. This approach is applied to other classes of appliance as follows. Heat pumps and electric vehicles are featured in the method description because of the particular challenge they present to distribution networks [5].

For cold appliances, the controller switches off the appliance for a single timeslot within the set q of timeslots for which $S_i > 0$, chosen with a probability:

$$P_{ci} = \frac{S_i}{\sum_{j \in q} S_j} \quad (3)$$

It is assumed that the appliance is a smart device which adjusts its temperature thresholds to redistribute the cooling load from the gapped timeslot uniformly over the remaining 47 timeslots. The timing of hot water heating by heat pump is determined by evaluating the attractiveness A_h of the default timeslot with heat demand H_d and each prior timeslot as:

¹ In this description of the home control unit algorithms a subscript notation is employed to identify attractiveness and probability variables comprising a first subscript indicating the appliance type (w – wet, c – cold, etc.) and a second subscript identifying the timeslot.

$$A_{hi} = \frac{(S_{\max} - S_i)H_d}{H_i} \quad (4)$$

taking into account cooling losses for earlier timeslots resulting in heat demand H_i . The calculation of probability over the q possible earlier timeslots is similar to (3):

$$P_{hi} = \frac{A_{hi}}{\sum_{j=1}^q A_{hj}} \quad (5)$$

Space heating by heat pump is controlled such that a gap in heating demand is introduced when $S_i > 0$ for one or two timeslots such that the fall in room temperature T_{ro} is less than 0.5°C . The heat pump is regulated by the control unit using the methods described in [18] such that the gap either provides energy saving or recovery of the heat input is spread over other timeslots so that there is no surge in demand after the gap. The calculation of probability is the same as (3) but extended to 2 timeslots where appropriate.

Charging of electric vehicles from the domestic supply follows the same concept. The home control unit learns the usage pattern of a vehicle under control so that when it is reconnected after use the control unit has a prediction of both the amount of charge Q_{tot} required to restore the battery to full capacity and the next time at which the vehicle will be needed. It then allocates this expected charge Q_{tot} over the r timeslots in which charging must be completed in proportion to the attractiveness $A_{vi} = (S_{\max} - S_i)$ of each timeslot. The charge Q_i in each timeslot is given by:

$$Q_i = \frac{Q_{tot}A_{vi}}{\sum_{j=1}^r A_{vj}} \quad (6)$$

3.2 Aggregator's model of demand response

Given a customer base of consumers equipped with control units as described above, it can be seen that the total demand in each timeslot seen by the aggregator will vary approximately proportionately to S_i but with different constants of proportionality k_i for each timeslot and for positive and negative values of S_i . These variations in proportionality arise from the differing constraints applied to each class of appliance across the consumer population. So for each of the 48 timeslots the change in demand stimulated by S can be modelled with two equations of the form:

$$\delta B_i = S_i k_i B_i + c_i \quad (7)$$

where k_i and c_i each have one of 2 values depending on whether S_i is positive or negative. Figure 3 shows an example of these 2 equations for a single timeslot obtained from a simulation of this concept.

Consequently if the aggregator can determine these 96 equations predicting the response to S it can construct S to achieve a range of objectives which are explored in the next section. It is assumed that the aggregator is able to predict the 48 value baseline demand vector B that is expected for any day, using established records and techniques such as neural network forecasting [19]. Then by transmitting different experimental constructions of S over a suitable period and recording the resulting difference δB_i in demand in each timeslot relative to the known baseline, the aggregator can construct a sufficient number of equations of the form given in (7) to solve for the 96 values of k_i and c_i . The demand D arising from any S can then be predicted using:

$$D_i = B_i(1 + S_i k_i) + c_i \quad (8)$$

where k_i and c_i are taken from the k and c vectors appropriate to the sign of S_i . Once in operation this model can be updated continuously from metering records. Any variation indicated in c_i must

be scrutinised to see if it is better considered as an incremental update to the baseline B_i as c_i values should represent effects of demand displacement. In practice the aggregator may wish to transmit a null signal to a proportion of consumers, either to sense the baseline demand directly, or to achieve a particular shape of aggregate demand. On days when demand is likely to be exceptional such as Christmas and New Year, the aggregator will have to expect the 96 equation model to be less accurate than usual and deploy the signal accordingly. Over time the aggregator will learn the weekly and seasonal variations in these parameters and adjust them accordingly.

The aggregator can also use the signal to convey a time-dependent price directly to the consumer to stimulate a demand response by human agency, such as timeshifting wet appliances without an automatic control capability or turning off unused entertainment appliances. This could be done by converting S in the control unit into a displayed price from a tariff rate vector \mathbf{R} with a relationship of the form $R_i = aS_i + b$. Alternatively price variations could be triggered by threshold values of S . As long as price changes are aligned in time and logically consistent with S then their effect on aggregate demand, which is likely to be fairly small as indicated by [6] and [7], can be represented simply as contributions to k and c in the response model.

4. Simulation of scheme

4.1 Domestic demand model

To evaluate the properties of this approach for shaping domestic demand, a model has been constructed which simulates the electricity consumption of a population of households equipped with a control unit and appliances that respond to a signal as specified in the previous section. In order to obtain realistic demands for space and water heating, the occupancy and thermal properties of these households are assigned as shown in Table 1. For each class of appliance, a default daily demand profile is created as summarised in Table 2. This profile provides the

default probability that the appliance is operating in any given timeslot. The demand when operating is calibrated to match either the average demand for this class of appliance from national statistics, or the actual thermal load of the simulated household as provided by the attributes from Table 1. The electrical load W_s for space heating is given in kW by:

$$W_s = \frac{L(T_r - T_e)}{p} \quad (9)$$

where p is the coefficient of performance (CoP) of the heat pump, L is the loss rate for the dwelling, T_e is the external ambient temperature taken from a test reference year [20], and T_r is the reference temperature (15.5 °C) which space heating must achieve, the remaining heat input for higher room temperatures being provided by appliance use and occupancy. This simplifying assumption results in space heating demand being directly dependent on external temperature, but this is typical of UK domestic heat pump demand where heat pumps are operated continuously to minimise heat emitter temperature and running costs, as illustrated by field measurements [21]. The change in room temperature T_{ro} when heating is ceased for one or two timeslots is calculated using a single node model of the building characteristics:

$$(T_{ro} - T_e) = \tau \frac{dT_{ro}}{dt} \quad (10)$$

$$\tau = \frac{C}{L}$$

where C is the thermal mass and L the loss rate from Table 1.

For water heating the electrical energy E_h required is calculated in kWh as:

$$E_h = VT_l \left(\frac{4.18}{3600p} \right) \quad (11)$$

where V is the hot water volume required by the household and T_l is the temperature lift for which a value of 40°C is taken from [21]. Cooling losses in hot water storage are calculated using a temperature fall of 0.4°C per hour since this is the specified performance for hot water storage vessels in UK building regulations [22]. COPs of 2.5 and 2.2 for space heating and water heating respectively are employed, based on Energy Saving Trust findings for UK installations [23]. These COPs are most typical of ground source heat pump installations, but some air-source installations documented in [23] achieve this level. Where Table 2 specifies “simulated aggregate profile” as the source of the daily profile for an appliance class, this refers to a fine-grained model of domestic appliance demand produced by Stokes [24, 25] which was employed to generate realistic inputs to this simulation.

The assumptions for simulation of electric vehicle use are given in Table 3. Baseline electricity use is determined by the profile of arrival times at home (shown in Figure 2) and trip distances derived from survey data [28] resulting in a distribution of charge requirements that are satisfied at a 3kW rate (constrained by the UK 13 Amp household socket).

Table 1. Household properties

Attribute	Values	Notes
Occupancy	Between 1 and 6 with a mean of 2.4	Distributed according to national statistics for 2010 [26].
Hot water use	$46+26n$ litres per day where n is occupancy.	Energy Saving Trust usage model for UK households [21]
Thermal loss rate of dwelling	Distributed between 0.05 and 0.4 $\text{kW}/^\circ\text{C}$ with a mean of 0.225	Consistent with national building energy ratings [27] and energy use [4]
Thermal mass of dwelling	Distributed between 5 and 20 $\text{kWh}/^\circ\text{C}$ with a mean of 12.5	Consistent with thermal loss rates and conventional UK construction

Table 2. Appliance default consumption profiles

Appliance type	Profile	Average demand (kWh/day)	Notes
Cold (fridge, freezer)	Flat	1.48	Average given total UK consumption in 2010 of 14TWh [4]
Wet (washing machine, dishwasher)	Simulated aggregate profile	1.52	Average given total UK consumption in 2010 of 14.4 TWh [4]
Heat pump (providing hot water)	Average hot water use profile	4.4	Profile from EST study [21]
Heat pump (providing space heating)	Determined by external temperature	29.5	Temperature profile from CIBSE TRY [20], average demand for a cold January day
Non-controllable (cooking, lighting, entertainment, computers)	Simulated aggregate profile	5.82	Average given total UK consumption in 2010 of 55 TWh [4]

Table 3. Electric vehicles

Attribute	Values	Notes
Vehicle population	A single car in 25% of households, of which 50% are battery-only, 50% plug-in hybrid.	Based on 75% of households with a car [28] and 28% of cars are EVs.
Battery capacity	24kWh for battery only EV, 16kWh for plug-in hybrid	Nissan Leaf [29], Vauxhall Ampera[30]
Energy consumption	0.17 kWh/km (battery-only), 0.2 kWh/km (plug-in hybrid)	From manufacturers' running cost specifications
Frequency of use	0.74 journeys/day, return time distributed as indicated in Fig.1	Derived from National Transport Survey 2010 [28]
Journey distance	Average 27km, Poisson distributed between minimum of 4 and maximum bounded by battery capacity.	Weekday modelled, distance from [28]
Recharge time	Distributed between minimum and 07:00 next day	Models constraint applied by user




Figure 2. Distribution of electric vehicle return times

For classes of appliance not amenable to management by the control unit (cooking, lighting, home entertainment and computing), an elastic response resulting from manual action is assumed, motivated by a display of the time of use tariff and real time consumption. Based on the elasticities of between -0.02 and -0.1 found by Faruqui and Sergici [7] this is modelled as a maximum response of -5% when $S=1$ and scaled proportionately for lower positive values. A lower response (maximum 2% when $S = -1$) is employed for negative values of S since while a high price may discourage some use of lighting or entertainment, a lower than average price is unlikely to attract additional use in low cost timeslots. However it is conceivable that some kitchen activities such as baking might be manually time shifted to cheaper timeslots.

The models of each household are able to return a baseline profile of electricity demand for each day, or the profile that results from applying the proportionate demand functionality described in 3.1, or the profile from an alternative response to S that results from the control unit seeking to minimise cost. In this latter mode the control unit is in effect executing similar behaviour to that described in references [8],[9],[10],[11], and [12]. This allows direct

comparison of proportionate and cost-minimising approaches. When the proportionate mode is in operation the simulation begins by allowing the baseline demand to emerge from the population of households, then a training sequence of values for S is transmitted by the simulated aggregator allowing the linear response model represented by equation (7) to be deduced. Finally the aggregator transmits a signal optimised to elicit a desired response depending on the scenario and the actual result is captured. Figure 3 illustrates the training response for a single i th timeslot by plotting $S_i B_i$ on the x- axis and the resulting δB_i on the y-axis for different values of S_i . The slopes of the regression lines either side of the y axis and their y-axis intercepts provide k_i and c_i values for positive and negative values of S . The timeslot shown in Figure 3 is 36 i.e. in the early evening where the baseline demand would be near a peak. The negative intercept for δB when $S_i = 0$ reflects the useful tendency noted in 3.1 for this method to spread demand from peaks to other timeslots which are equally or more attractive.




Figure 3. Linear response of aggregate demand to S for an example timeslot (36)

4.2 Results

Most of the results presented below share a common scenario of 1000 households with space and water heating by heat pumps, and appliance and electrical vehicle use, as specified in Tables 1-3. A winter day with overnight temperatures falling to $-1.4\text{ }^{\circ}\text{C}$ is simulated in all cases as this is likely to represent a worst-case load on local distribution networks. A typical resulting baseline demand profile from these households is shown in Figure 4 including the contributions from the various appliance types – the miscellaneous type covers all those appliances such as entertainment devices that are only amenable to manual control. To demonstrate the inherent demand flattening propensities of this scheme, a comparison is provided in Figure 5 of the total demand profiles resulting from proportionate and cost-minimising responses in the control unit when the same population is presented with a price signal that follows the shape of the national aggregate demand on a winter day. The peak-mean ratios of the baseline, proportionate, and cost-minimising responses respectively are 1.4, 1.3, and 2.2. The demand-flattening potential is fully exploited when the aggregator uses the demand response model to generate a signal optimised to induce a flat response, with the objective function from equations (1) and (8):

$$\text{Minimise: } \frac{D_{\max}}{D} \quad (12)$$

$$\text{where } D_i = B_i(1 + S_i k_i) + c_i ;$$

$$-1 \leq S_i \leq 1 \quad \text{for } i = 1:48;$$

$$\sum_{i=1}^{48} S_i = 0 .$$

The simulation employs the “active set” method [31] to solve this non linear optimisation problem for S . The response of the consumer population with a baseline demand as indicated in Figure 4 to a signal S determined in this way is shown Figure 6, which has a peak-mean ratio of

1.09. The residual variations in demand arise from the imperfections in the demand model held by the aggregator and the various constraints on electricity use applied by consumers.




Figure 4. Baseline winter demand from consumer model – 1000 households

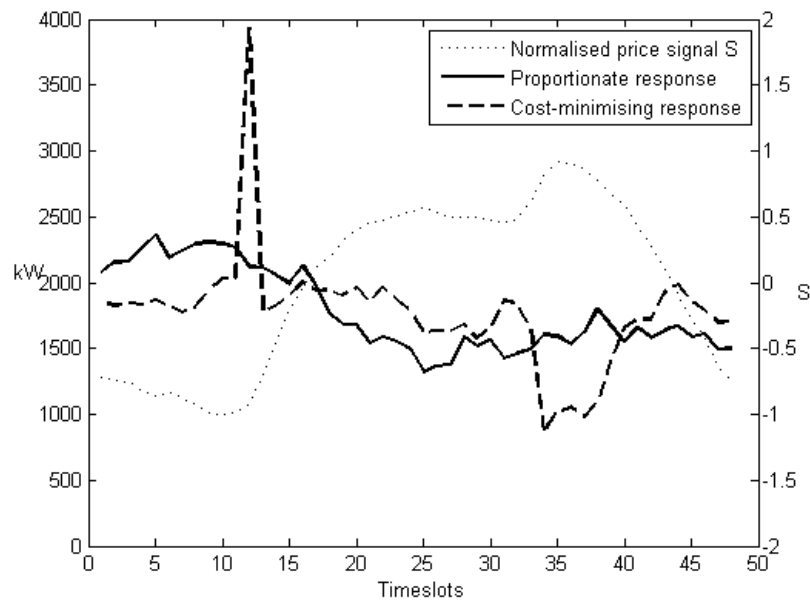


Figure 5. Comparison of proportionate and cost-minimising control algorithms




Figure 6. Demand response to a signal optimised for flattening

The aggregator can employ a suitably optimised S to obtain any desired shape of demand within the flexibility limits determined by consumer preferences and needs and the constraints in (12). Figure 7 shows the upper and lower limits to the demand profiles that can be induced in this scenario, with the baseline for comparison. The upper limit is below the baseline peak because the control unit algorithms result in minimum levels of demand displacement from a non-zero S as indicated by the c values in the aggregator model (7).

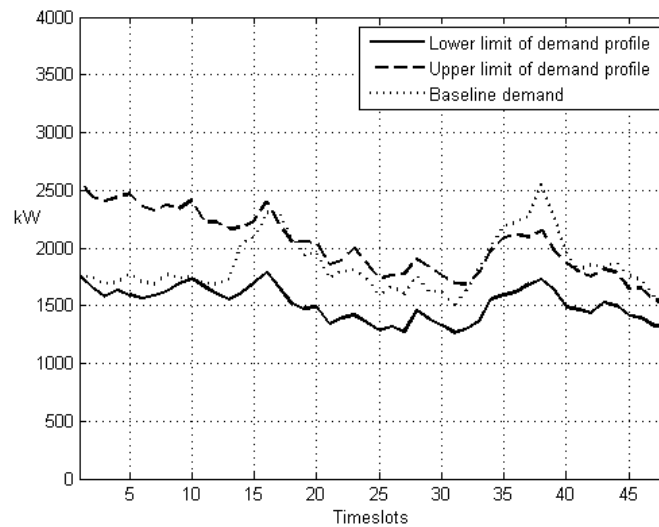


Figure 7. Upper and lower limits to demand response

This flexibility allows demand to be adjusted to match the varying availability of renewable generation by replacing the objective function in (12) with one that minimises the difference between demand and a desired target profile. Figure 8 portrays a scenario in which some of the 1000 households modelled have photovoltaic (PV) panels installed with a total capacity of 350kWp and the weather forecast is for a sunny but cold winter day. Figure 8 shows actual demand, rather than net demand when PV output is subtracted, so that the induced variation in demand can be seen. The aggregator is able to increase demand between timeslots 20 and 30 to the limit indicated in Figure 7, amounting to an additional 450kWh in that interval, with a signal optimised against a target profile shaped to reflect PV output. This is only 1.2% of the electricity consumed during the day but at the margin it may be valuable to an aggregator who is obliged to pay feed-in tariffs to customers with PV panels and provides a useful service to those customers. Alternatively the aggregator might be a community energy co-operative with the PV panels as part of its generating assets, which need to be exploited as effectively as possible. The signal shown is that produced by the optimisation algorithm to achieve this result.

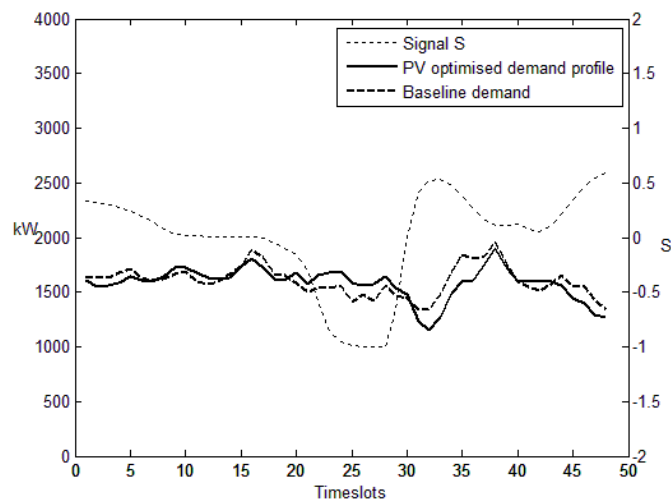


Figure 8. Demand optimised for PV generation availability




Figure 9. Demand optimised for overnight wind generation availability

In Figure 9 it is envisaged that the aggregator has access to a surge of wind generation overnight which is expected to reduce later in the day. A signal optimised to a target profile lifting demand between timeslot 0 and timeslot 15 attracts an additional 3300 kWh of demand into the early hours amounting to 7.7% of the demand for the day. Had the wind surge been expected later in the day the aggregator could simply have sent a null signal allowing the baseline demand peak to occur.

4.3 Provision of despatchable reserve capacity to support wind generation

To see how this scheme could contribute to the despatchable reserve capacity identified in [3] as necessary to support expansion of UK wind generation, a national scale simulation was performed of a 2020 scenario assuming that the central estimates of wind generation capacity (26GW) and heat pump rollout (1.2 million homes) in the UK Renewable Energy Roadmap [1] are achieved. It is also assumed that half the existing population of 1.7 million homes heated by electricity [32] using thermal storage radiators are brought into the scheme. To reflect the fact that some of the heat pumps are likely to be operated on a timed basis similar to common practice for gas boilers, the expected air source proportion (400,000) are modelled with timed morning and evening operation typical for households unoccupied during the day. This results in a more

“peaky” baseline demand pattern than shown in Figure 4. The national wind generator output was simulated using a similar technique to that used in [9]. The output profile is generated from 1500 simulated wind farms, with capacities appropriately scaled to give a total nominal capacity of 26GW. Weather variation across the UK is simulated by using a wind speed time series for a single site [20] and offsetting it by a randomly selected number of half hour timeslots up to +/- 1 day before it is experienced by each wind farm. This approximates the passing of weather fronts over the UK, with a particular wind speed occurring at one edge of the country a maximum of 2 days before it arrives at the other.

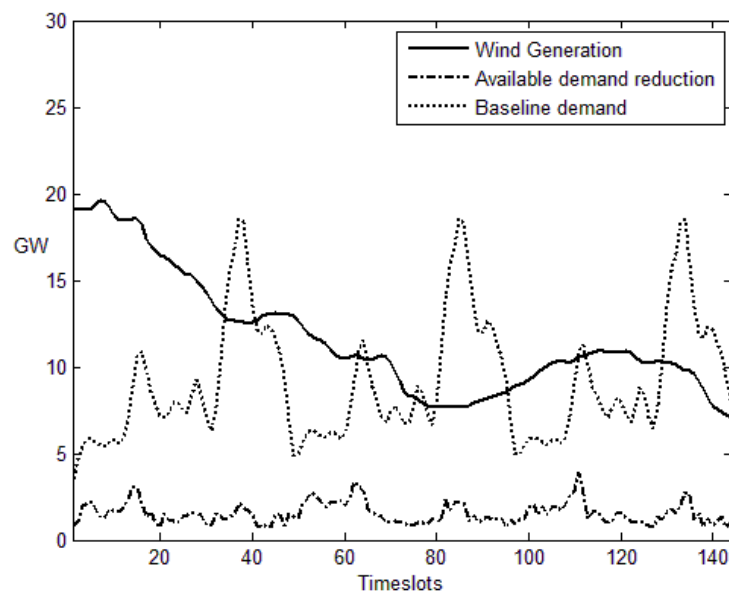


Figure 10. Use of demand response as dispatchable reserve capacity

Three days in winter from this simulation are shown in Figure 10 which feature a particularly large change in wind generation from a capacity factor of 75% at timeslot 7 to 27% at timeslot 144. The demand reduction shown as available at each timeslot is calculated by subtracting the lower limit of demand (as shown in Figure 7) from the baseline demand. This is equivalent to dispatchable reserve capacity given that the displaced demand is shifted to a time when it can be supplied more efficiently, such as around timeslot 30 where a dip in baseline demand can be seen in Figure 10. In this scenario this reserve capacity varying from 0.7-2.8GW

would probably be deployed to reduce the three early evening peaks by about 2.2 GW. The total despatchable generation needed to manage both wind and demand variability over these 3 days is 28GW so this scheme would contribute about 8% of this total and represent about 78% of the additional despatchable generation specified in [3] as necessary to support the wind capacity.

5. Discussion

The key advantage of using demand response as despatchable reserve capacity is that it avoids the additional carbon emissions that would otherwise occur from operating fossil fuelled plant that is low in the merit order and only brought into play to meet demand peaks. It should also be cheap to operate, since the infrastructure required is largely being provided under the national programme to roll out smart metering. The additional cost of the smart control unit should be amply repaid by the large savings in distribution reinforcement costs identified in [5], so the reserve capacity provided to support wind generation should have very low operational cost. The only disadvantage is that despatching has to be performed about a day ahead, but given the reasonable reliability of short term weather forecasts this should not be an issue. If the policy for renewable energy in [1] is implemented, as electrification of heating and transport proceeds this reserve capacity will grow in parallel with growth in renewable generation.

The winter scenario used to illustrate this demand response scheme was chosen because it is a critical case for the electricity supply industry but a greater level of response would be available under other scenarios. In particular space heating gaps were limited to 1 hour to minimise impact on perceived comfort. In practice many households, depending on the climatic conditions, would be able to tolerate longer or multiple gaps providing much more demand flexibility. The displacement of water heating load was also limited by the condition that it could only be moved to a time earlier in the same day. If S was signalled for 2 days ahead, with the provision to update the second day, attractive timeslots late in the first day could be exploited to meet morning hot water demands in the second day. More frequent updates could also be

performed to suit the dynamics of a particular pool or other market structure, but the resulting response will be constrained by the natural cycles in consumer demand such as commuting to work in an electric car. The impact of electric vehicles appears modest because of the scale of space heating, but at other times of year, and with vehicle-to-grid flexibility engaged, they could have a more substantial and evidently valuable role.

Despite these limitations the simulation results show that this method for inducing demand response is capable of both moving demand in time to coincide with renewable electricity generation and reducing demand peaks. Figure 9 shows both objectives being achieved simultaneously. It also has several other useful properties:

- It is completely scalable in that the signal can be directed at a few households or millions. If the aggregator builds the models expressed in equation (7) at the lowest level of decomposition of its consumer population at which discrete profile shaping is likely to be required, then larger groupings of consumers up to the total population can be modelled by summing the linear models from the lowest level.
- For many useful demand profiles, such as the flattened profile shown in Figure 6, the amount of electricity consumed could be reduced through user action to switch off some loads when the signal is interpreted and displayed as a higher price.
- A single home control device can manage all appliance types and electric vehicles using a common underlying concept. There is no need for a specialised electric vehicle charging controller as proposed by some authors [33] as the concept is readily adaptable to “vehicle to grid” schemes by allowing limited and proportionate export of energy from the vehicle battery in timeslots where S is positive. It is similarly applicable to despatchable micro generation such as micro combined heat and power plants. This simplification and commonality will aid user understanding and acceptance of the principle of demand response.

- Because the PAR of typical demand profiles is always reduced the gain in this element of the demand response feedback loop is less than 1 so stability is enhanced in wholesale market pricing and total demand.

A critical issue that this scheme may help to resolve is the relationship between aggregators and distribution network operators (DNOs). In many countries an “arm’s length” relationship is enforced between them by regulation because the former compete with each other to supply consumers while the latter are natural monopolies. However both have an interest in shaping demand to meet the two objectives outlined in the introduction. A process that balanced their interests might take the following form:

- The DNO specifies a 48-value maximum power profile for each connection point using rules defined by the regulator and provides this to the aggregator holding the supply relationship for that connection.
- The aggregator assesses the risk of each consumer exceeding their maximum power profile from their historic metering data and allocates them to a group to receive an appropriate signal depending on the risk. Thus consumers with little headroom would be sent a signal prioritised for flattening such as that simulated in Figure 6, whereas consumers connected with more network capacity could be given a signal attracting demand at times when renewable generation is available as shown in Figures 8 and 9.
- A regulated audit process would identify consumers who exceed their profile from metering data and incentivise the aggregator to minimise occurrences. Where the demand is being flattened as far as possible through induced response and the profile limits are still being exceeded then that would be evidence to justify network reinforcement or a physical tripping limit on consumption.

This process would ensure that the DNO's network capacity constraints are respected without any need for the DNO to interact directly with consumers on a routine basis or have access to their individual metering data (but DNOs would have metering totals for a feeder or geographical segment). The aggregator would retain the primary relationship with the consumer and be responsible for the confidentiality and privacy issues in the processing of metering data.

6. Conclusion

This investigation has shown that electricity demand from domestic consumers equipped with half-hour metering can be shaped to address critical pressures on the electricity supply industry through the adoption of an automated demand response mechanism implemented by a "smart home" control unit that randomises the response of an individual consumer with a bias depending on the signal provided by their electricity supplier. Such a scheme if implemented on a large scale would allow the benefits of demand response – reduced need for peaking plant, better exploitation of renewables, and more efficient use of distribution network assets - to be realised without risk to system and market stability.

Acknowledgement

The authors would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for providing the financial support for this study as part of the CASCADE project (EP/GO59969/1).

Nomenclature

A	Attractiveness vector derived from signal S
a	Slope of linear function for converting S to a price
B	Baseline electricity demand in day (48 values)
b	Intercept of linear function for converting S to a price
C	Thermal capacity of building (kWh/°C)
c_i	Intercept of aggregator's model for i th timeslot
D	Electricity demand in day (48 values) including response to S
E	Electrical energy (kWh)
i, j, q, r	Index to a half-hour timeslot
H_d, H_i	Heat required in the default timeslot and i th earlier timeslot
k_i	Slope of aggregator's model for i th timeslot
L	Heat loss rate of building (W/°C)
P	Probability vector of electricity consumption in a time interval
p	Coefficient of Performance of heat pump
Q_{tot}, Q_i	Total charge required by electric vehicle, charge in i th timeslot
R	Price vector derived from signal S
S	Signal vector (48 value) influencing electricity demand
T_e, T_r, T_{ro}, T_l	External, reference, room, and lift temperature (°C)
t	Time, in timeslots (integer) or elapsed
τ	Building time constant
V	Volume (litres)
W	Electrical load (kW)

References

- [1] Department of Energy and Climate Change: ‘The UK Renewable Energy Roadmap’, <http://www.decc.gov.uk/assets/decc/11/meeting-energy-demand/renewable-energy/2167-uk-renewable-energy-roadmap.pdf> accessed February 2013
- [2] Department of Energy and Climate Change: ‘2050 pathways analysis’, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/68821/2050-pathways-analysis-response-pt1.pdf accessed February 2013
- [3] Gross, R., Heptonstall, P., Anderson, D., Green, T., Leach, M., and Skea, J.: ‘The costs and impacts of intermittency: an assessment of the evidence on the costs and impacts of intermittent generation on the British electricity network’ (UK Energy Research Centre, 2006).
- [4] Department of Energy and Climate Change: ‘Energy consumption in the UK – domestic data tables’, <http://www.decc.gov.uk/en/content/cms/statistics/publications/ecuk/ecuk.aspx> accessed February 2013
- [5] Pudjianto, D., Djapic, P., Aunedi, M., Gan, C. K., Strbac, G., Huang, S., Infield, D.: ‘Smart control for minimizing distribution network reinforcement cost due to electrification’, *Energy Policy*, 2013, 52, pp. 76–84
- [6] Hargreaves, T., Nye, M., and Burgess, J.: ‘Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term’, *Energy Policy*, 2012, doi.org/10.1016/j.enpol.2012.03.027
- [7] Faruqui, A., and Sergici, S.: ‘Household response to dynamic pricing of electricity: a survey of 15 experiments’, *J Regul Econ*, 2010, 38, pp. 193-225.
- [8] Boait, P., Rylatt, R.M., and Wright, A.: ‘Exergy-based control of electricity demand and microgeneration’, *Applied Energy*, 2007, 84, pp. 239-253.
- [9] Roscoe, A., and Ault, G.: ‘Supporting high penetrations of renewable generation via implementation of real-time electricity pricing and demand response’, *IET Renew. Power Gener.*, 2010, 4(4), pp. 369-382.
- [10] Papadaskalopoulos, D., and Strbac, G.: ‘Decentralised agent-based participation of load appliances in electricity pool markets’, *Proc 21st Int. Conf. on Electricity Distribution (CIRED)* Frankfurt, June 2011, paper 1049

- [11] Mohensian-Rad, A., and Leon-Garcia, A.: ‘Optimal residential load control with price prediction in real-time electricity pricing environments’, *IEEE Trans. Smart Grid*, 2010, 1(2), pp. 120-132.
- [12] Ramchurn, S.D., Vytelingum, P., Rogers, A., and Jennings, N.R.: ‘Agent-Based Control for Decentralised Demand Side Management in the Smart Grid’, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Taipei, May 2011, pp. 5-12.
- [13] Rastegar, M., Fotuhi-Firuzabad, M., and Aminifar, F.: ‘Load commitment in a smart home’, *Applied Energy*, 2012, doi:10.1016/j.apenergy.2012.01.056.
- [14] Giordano, V., and Fulli, G.: ‘A business case for Smart Grid technologies: A systemic perspective’, *Energy Policy*, 2012, 40, pp. 252-259.
- [15] Vaze, P., and Tindale, S.: ‘Repowering communities – small scale solutions to large scale energy problems’ (Earthscan, 2011)
- [16] Strbac, G.: ‘Demand side management: benefits and challenges’, *Energy Policy*, 2008, 36, pp. 4419–4426
- [17] Glass, J., Dainty, A.R., and Gibb, A.: ‘New build: materials, techniques, skills and innovation’, *Energy Policy*, 2008, 36, pp. 4534–4538.
- [18] Boait, P.J., Fan, D., and Stafford, A.: ‘Performance and control of domestic ground-source heat pumps in retrofit installations’, *Energy and Buildings*, 2011, 43, pp. 1968-1976.
- [19] Hippert, H.S., Pedreira, C.E., and Souza, R.C.: ‘Neural networks for short-term load forecasting: a review and evaluation’, *IEEE Trans. on Power Systems*, 2001, 16 (1), pp. 44 – 55.
- [20] CIBSE: ‘Current CIBSE TRY/DSY Hourly Weather Data Set - Birmingham 2010’, <http://www.cibseknowledgeportal.co.uk/weather-data> accessed February 2013.
- [21] Boait, P., and Stafford, A.: ‘Electrical load characteristics of domestic heat pumps and scope for demand side management’, *Proc 21st Int. Conf. on Electricity Distribution (CIRED)* Frankfurt, June 2011, paper 0125
- [21] Energy Saving Trust: ‘Measurement of Domestic Hot Water Consumption in Dwellings’, <http://www.bsria.co.uk/news/est-water/> accessed February 2013

- [22] Department for Communities and Local Government: ‘Domestic building services compliance guide’, 2010.
- [23] Energy Savings Trust: ‘Getting warmer: a field trial of heat pumps’(Energy Saving Trust, 2010).
- [24] Stokes, M.: ‘Removing barriers to embedded generation: a fine-grained load model to support low voltage network performance analysis’. PhD thesis, De Montfort University, 2005.
- [25] Stokes, M., Rylatt, M., and Lomas, K.: ‘A simple model of domestic lighting Demand’, *Energy and Buildings*, 2004, 36, pp. 103–116.
- [26] Office for National Statistics: ‘Families and households 2001-2011’, <http://www.ons.gov.uk/ons/publications/re-reference-tables.html?edition=tcM%3A77-248983> accessed February 2013
- [27] Department for Communities and Local Government: ‘English Housing Survey Headline Report 2010-11’, <http://www.communities.gov.uk/publications/corporate/statistics/ehs201011headlinereport> accessed February 2013
- [28] Department for Transport: ‘National Travel Survey 2010’, <http://www.dft.gov.uk/statistics/releases/national-travel-survey-2010> accessed February 2013
- [29] Nissan Motor GB Ltd: ‘Nissan Leaf specification 2012’, <http://www.nissan.co.uk/#vehicles/electric-vehicles/electric-leaf/leaf> accessed February 2013
- [30] General Motors UK Ltd: ‘Vauxhall Ampera specification 2012’, <http://www.vauxhall.co.uk/vehicles/vauxhall-range/cars/ampera/index.html> accessed February 2013
- [31] Gill, P. E., Murray, W., and Wright, M.H.: ‘Practical Optimization’ (Academic Press, 1981)
- [32] Department for Communities and Local Government: ‘English housing survey housing stock summary statistics tables (2009) pp. SST6.2’, <http://www.communities.gov.uk/documents/statistics/xls/1937429.xls> accessed February 2013
- [33] Rotering, N, and Ilic, M.: ‘Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets’, *IEEE Trans. Power Systems*, 2011, 26(3) pp. 1021-1029.