

Consumer Targeting in Residential Demand Response Programmes

James C. Holyhead
Agents, Interaction and
Complexity Group
Electronics and Computer
Science
University of Southampton, UK
jch2g12@ecs.soton.ac.uk

Sarvapali D. Ramchurn
Agents, Interaction and
Complexity Group
Electronics and Computer
Science
University of Southampton, UK
sdr@ecs.soton.ac.uk

Alex Rogers
Agents, Interaction and
Complexity Group
Electronics and Computer
Science
University of Southampton, UK
acr@ecs.soton.ac.uk

ABSTRACT

Demand response refers to a family of techniques that are available to electricity suppliers to aid with balancing supply and demand, typically by calling on consumers of electricity to reduce consumption during periods of high demand. In this paper we propose a novel approach to residential demand response, in which incentives are targeted at the subset of consumers who are both relevant (likely to use shiftable appliances, such as washing machines and dishwashers during peak hours) and willing to reduce (likely to react positively to a reduction request from their electricity supplier). To this end, we present a mixed integer programming solution that finds the optimal subset of consumers to target with incentives. We show that our solution is capable of significantly reducing supplier costs and smoothing peaks in electricity demand by targeting only a subset of the consumer pool.

1. INTRODUCTION

Aging electricity distribution networks, increasing demand and ambitious CO_2 reduction targets have caused many developed countries to consider how to best upgrade their electrical networks to meet the demands of the 21st century [12]. The ‘smart grid’, a network that facilitates a two way flow of both electricity and information, is an emerging solution to this problem of next generation electricity distribution [13]. The introduction of smart grid technologies such as the smart meter have made it possible for the suppliers of electricity to communicate with consumers in real-time. These messages may be in the form of pricing signals, which are used to incentivise consumers to alter their electricity consumption as part of so-called ‘demand response’. Demand Response is a family of approaches that call upon electricity consumers to curtail or to reschedule their electricity usage in response to requests from their electricity supplier [14] or from the grid operator. These consumers are typically offered incentives for compliance with these requests or

penalties for failing to comply [2]. These requests are often issued at times of particularly high demand or as a response to shortfalls in generation caused by equipment failures on the grid. Customer responses to these requests alter the shape of the aggregate demand profile, often by reducing or smoothing peaks in demand.

Existing demand response approaches targeted at residential consumers have been focussed on tariff based methods [11, 6]. These schemes make use of differential pricing, where the cost per unit of electricity varies periodically throughout the day in an attempt to nudge consumers into altering their consumption behaviour. For example, during periods of high load, the price of electricity will be increased to dissuade consumers from activating non-essential appliances. One subcategory of differential pricing is ‘real-time pricing’ (RTP). In an RTP scheme, the price of electricity is varied on an hourly or half hourly frequency with the prices being announced on a day ahead or hour ahead basis [2]. Economists typically consider RTP to be the most efficient differential price mechanism for use in demand response programs [1]. However, we argue that such schemes place a high information burden on consumers, which leads to poor consumer retention rates. In 2005, Rocky Mountain Power in Utah, USA evaluated the use of their differential pricing based tariffs. They found that in opt-out schemes, up to 98% of participating consumers chose to leave the programme after the mandatory period had been completed [9]. This indicates that the participants of differential pricing schemes do not feel that the inconvenience of rescheduling their appliance usage is sufficiently compensated for by the cost savings they earn as participants.

The impact of this additional information burden can be seen in the prevalence of ‘demand response fatigue’ in residential demand response trials [8]. This fatigue results in a progressive disengagement from the scheme over time – as the novelty of a new tariff wears off, consumers begin to ignore the price signals and their consumption behaviour reverts to the behaviour expressed under static pricing, which often leads to a net increase in the costs incurred by the consumer when compared to static pricing. If a scheme cannot retain its participants, then it is not sustainable.

Several solutions have been proposed to the problem of residential demand response. Xiao et. al [17] model the problem

as a task scheduling problem, but the authors find that their solution is intractable for non-trivial numbers of consumers. Ramchurn et.al. [11] aim to find a trade-off between the cost savings that a consumer can make by rescheduling appliances against a comfort cost¹ assigned to the rescheduling of appliances. However, they do not capture the full range of consumer behaviour in their model. In Chandan et. al. [3] the authors rank consumers according to the flexibility of their appliance usage in an attempt to target those consumers who would be the least inconvenienced by reduction requests, but the complexity of their algorithm means the number of consumers they can consider is limited.

To address these shortcomings, we propose a novel approach to demand response, where incentives are targeted at those electricity consumers who are both ‘relevant’ (likely to use shiftable appliances, such as washing machines and tumble dryers during peak hours) and ‘flexible’ (likely to react positively to an incentivised reduction request from their electricity supplier). We will show that the targeting of incentives to individual consumers reduces the costs incurred by the supplier, reduces peaks in demand. These individual requests may also reduce the information burden on the consumer, which in turn may lead to greater consumer retention rates.

We focus our attention on the shifting of ‘shiftable static loads’, such as dishwashers, washing machines and tumble dryers.¹ The shiftable static loads we concern ourselves with require direct consumer interaction; a dishwasher cannot be used before the consumer has filled it with dirty dishes and similarly, a washing machine cannot be used until the consumer has loaded it with dirty laundry. Further, there is evidence that consumers prefer to schedule the use of these appliances manually, even when a system is available to automate the process [5].

Thus, in this paper, we advance the state of the art in the following ways:

- We present a novel model of customer behaviour that captures the differences in how consumers respond to reduction requests, both in terms of their *willingness* to shift consumption, but also *how* the consumer will shift consumption (pre-emption or deferral).
- We propose a novel mixed integer programming solution to select the optimal subset of consumers that should be targeted with demand reduction requests and for which time period those reduction requests should be issued.
- We evaluate our solution using a dataset of real-world consumption data and show that our targeting algorithm and consumer response model lead to greater reductions in peak demand and lower supplier costs when compared to current models of consumer behaviour. We also show that our solution only requires participation by a portion of the consumer pool, reducing the

¹We omit any discussion of thermal loads or electric vehicle charging, as the scheduling of these loads may be automated by an intelligent agent without impacting on user comfort.

impact of ‘demand response fatigue’, which could lead to improved retention rates.

The rest of this paper is organised as follows. In Section 2 we discuss related work in the field of residential demand response. Section 3 presents a formalisation of the demand response problem. Section 4 describes our novel consumer response model and our approach to finding the optimal subset of consumers to target with demand reduction requests. Section 5 provides an empirical evaluation of our model and our consumer targeting approach and Section 6 provides concluding remarks and a discussion of future work.

2. BACKGROUND

The problem of rescheduling appliance operation in order to effect peak reduction has been well studied. In this section we outline some of the different solutions that have been proposed to the problem of residential demand response.

Xiao et. al. [17] addressed the optimisation problem as a task scheduling problem where each load request is given an earliest start time and a deadline for completion. Their solution seeks to find the optimal mapping of appliance usage requests to timeslots in order to minimise the cost within an individual timeslot. The authors note that this task scheduling formulation is NP-hard and therefore intractable for large numbers of demand requests. By contrast, we show that the solution proposed in this paper is capable of scaling to communities of tens of thousands of consumers.

This in turn contrasts with Ramchurn et. al. [11], who propose a solution using a linear programming formulation to the deferrable load scheduling problem and include an explicit measure of consumer comfort. A scaling factor is used to denote a consumer’s preference for comfort or cost saving. This measure of consumer comfort acknowledges that asking a consumer to reschedule their appliance usage carries some non-financial penalty in terms of convenience, but the authors stops short of defining a probabilistic measure of a user’s willingness to shift. In addition, they scale this comfort cost by the size of the shift, acknowledging that consumers will often prefer to reschedule their consumption as little as possible, but this does not capture the full range of potential consumer responses to reduction requests. The consumer response model we propose in this paper allows for a much richer representation of consumer behaviour in demand response settings, which we show leads to greater reductions in peak reduction performance.

In Chandan et. al. [3], the appliance scheduling optimisation problem is formulated as a mixed integer non linear program (MINLP). Consumers are given a score that indicates how greatly they would be inconvenienced by a load shifting request. The consumers are ranked in order of their flexibility with the highest ranked being used as input for the MINLP solver, representing the consumers who would be least inconvenienced. The remaining consumers are discarded to reduce the complexity of the problem. Rather than ranking consumers in terms of their convenience, the solution we propose in this paper targets consumers whose expected behaviour leads to a reduction in the supplier’s expected costs. Additionally, we do not need to pre-select the

consumers we examine as our solution is capable of solving problems for tens of thousands of consumers.

Some commercial demand response schemes have tackled the problem of peak reduction by issuing reduction requests to their scheme participants for times of high demand [10]. We will benchmark the performance of our solution against this approach later in this paper.

None of these approaches take into consideration the full variety of potential consumer responses to a demand response event, whether it be triggered by a differential pricing scheme or other form of incentive. The unique socio-technical challenges posed by residential demand response programmes are illustrated in the ‘Agent B’ field study [5], where an intelligent booking system was deployed to help consumers make efficient use of their washing machines in the presence of an RTP tariff. The authors discovered that even when the system offered to automate the booking process by finding the optimal operating times, the study participants preferred to carry out the scheduling manually, using the booking system in an advisory capacity only.

In the following section we formalise the problem of residential demand response.

3. THE DEMAND RESPONSE PROBLEM

Electricity is an unusual commodity in that it is not feasible to store it in large quantities due to the high cost of storage devices [4]. This means that electricity must be consumed at the same time it is generated and the supply must always match (be in balance with) the demand to prevent system instability and possible collapse.

The majority of energy is traded through bilateral contracts between the generators and the suppliers, which are often made far in advance of the time of consumption. Typically about half of the total volume of electricity traded within Great Britain is traded ‘seasons ahead’ [16].

The nature of electricity demand means it is impossible for a supplier to accurately predict their demand requirements months or even years ahead of time. Electricity markets enable suppliers to adjust the amount of electricity they are contracted to purchase through energy exchanges, where the various electricity suppliers can buy and sell electricity from one another. As the day the electricity has been contracted for approaches, the supplier can fine-tune their position through these exchanges.

We denote the volume of electricity that a supplier has contracted for as $\mathbf{\Gamma} = \{\Gamma_1 \cdots \Gamma_T\}$ and the total electricity demand consumed as $\mathbf{B} = \{B_1 \cdots B_T\}$.

We define any consumption above $\mathbf{\Gamma}$ to be a deficit, $\mathbf{d} = \{d_1 \cdots d_T\}$, where the deficit at time t , d_t is defined as follows:

$$d_t = \max(0, B_t - \Gamma_t) \quad (1)$$

Similarly, any consumption below $\mathbf{\Gamma}$ is defined as a surplus, $\mathbf{s} = \{s_1 \cdots s_T\}$ where the surplus at time t , s_t is defined as follows:

$$s_t = \max(0, \Gamma_t - B_t) \quad (2)$$

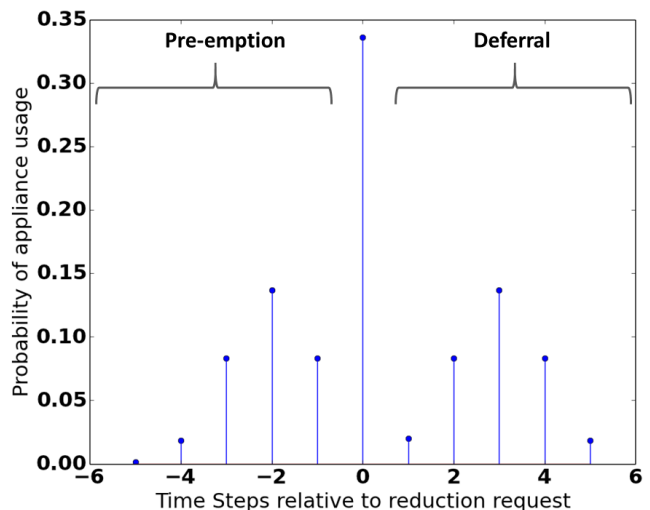


Figure 1: An example of a customer response profile. In this example we see the probability of the consumer ignoring the reduction request ($\tau = 0$) is approximately 0.35.

A deficit at time t , d_t , will incur a cost per kWh of ρ_{buy}^t where $\rho_{buy}^t > 0 \quad \forall t$ and a surplus at time t , s_t will be sold through the balancing market at a cost per kWh of \hat{p}_{sell}^t , where the penalty cost incurred by the supplier, ρ_{sell} is the difference between the price per kWh originally paid by the supplier for the electricity, ρ_t and \hat{p}_{sell}^t .

The prices set by the balancing mechanism, ρ_{buy} and \hat{p}_{sell} are set higher and lower than normal market rates respectively. This is to disincentivise suppliers from making use of the balancing mechanism as anything other than a last resort.

Given these market conditions, a supplier wishes to initiate a demand response event that leads to a reduced deficit or surplus and therefore a lower cost.² In the next section we present an optimal solution to this problem that identifies which consumers to target with demand reduction requests in order to minimise the expected cost to the supplier.

4. AN OPTIMAL MILP SOLUTION TO THE DEMAND RESPONSE PROBLEM

In this section we provide details of our model of consumer response behaviour and present an optimal solution to the problem of targeting incentives to consumers.

4.1 Consumer Response Model

For a supplier of electricity, the ideal outcome of a demand response scheme is a smooth profile (one with no peaks). Achieving this goal is more complicated than simply incentivising consumers to reduce their electricity consumption

²Although we use supplier cost minimisation as the motivating example in this paper, the same approach can be applied to any scenario where deviation from baseline $\mathbf{\Gamma}$ leads to some penalty being incurred. An alternative to a financial penalty may be increased CO_2 emissions.

during periods of high demand because large numbers of those consumers may activate their deferred appliances immediately after the reduction event has passed. This synchronisation of rescheduled appliances may lead to one or more secondary peaks that pose the same problems as the original peak [11]. The tendency for shifted loads to synchronise and form secondary peaks is known as the ‘pay-back effect’ [7]. A successful peak reduction programme must therefore account for the potential emergence of these secondary peaks. In this section we propose a model of consumer behaviour that takes into account how a user reschedules her consumption in the event a reduction request is issued. Rather than modelling a consumer’s willingness to shift as a single probability value, we model a user’s demand response behaviour by a probability mass distribution that describes the probability of when a shifted appliance activation has been rescheduled for. For example, a washing machine may be activated 2 hours after a demand response period (a deferral) or it may be activated 2 hours prior to the demand response period (a pre-emption). An example of such a consumer response distribution is shown in Figure 1. By accounting for the variety in consumer responses, we are able to simulate aggregate consumer behaviour that results in a smoother consumption profile.

We represent a user’s demand response behaviour by a probability mass distribution defined as:

$$p'_t(\tau) = P(z_{t+\tau} = 1) \quad (3)$$

where $z_{t+\tau} \in \{0, 1\}$ indicates whether the consumer reschedules their shiftable appliance usage for the time period $t + \tau$ given that they received a reduction request from their supplier for time t . We represent the length of the shift by $\tau \in \{-T, \dots, T\}$. In the case that $\tau = 0$, the consumer has chosen not to shift or curtail their consumption and $\tau < 0$ indicates that the consumer has pre-empted their consumption and $\tau > 0$ indicates a deferral of consumption. A consumer’s response distribution can be determined by analysing their behaviour during historical demand response events.³

We will use this probability distribution later when computing the consumer’s updated demand following a demand reduction request.

This approach differs from current practice by modelling a consumer’s behaviour in terms of *how* they shift their consumption. Existing approaches often assume that consumer’s will always respond to price signals/incentives when doing so would lower the price of their appliance usage [11] or only consider whether a consumer’s response in terms of accept/reject without considering how the rescheduled load may itself lead to unwanted system level behaviour, such as a secondary peak [3].

4.2 Optimal Consumer Selection

In what follows, we present an optimal solution to finding sets of consumer demand reduction recommendations that reduce the cost of the aggregate demand profile to the sup-

³The detail of the learning algorithm are beyond the scope of this paper and is left for future work.

plier. This solution takes the form of a mixed integer linear program.

In a demand response programme, the supplier aims to select a subset of consumers whose predicted behaviour during demand response events minimises the cost of the aggregate demand profile produced by its consumer’s demand requirements.

The total cost to the supplier of balancing their demand can therefore be calculated as follows:

$$C = \sum_{t=0}^T s_t \rho_{sell}^t + d_t \rho_{buy}^t \quad (4)$$

Note that only one of the two clauses are non-zero at each time step because it is not possible for both d_t and s_t to be greater than zero (the supplier cannot simultaneously be experiencing a surplus and a deficit), so at least one of the clauses will be cancelled out at each timestep. The supplier wishes to initiate a demand response event that results in an aggregate demand profile that costs less than C .

Given a set of consumers \mathcal{I} and a day split into T time slots, we define $\mathbf{B} = (B_1, B_2, \dots, B_T)$ to represent the baseline consumption profiles of the consumer’s within the supplier’s customer pool and we define C_B to be the cost of the aggregate consumption profile to the supplier (see equation 4). We now compute \mathbf{B}' , a $(T + 1) \times |\mathcal{I}|$ matrix where $B'_{t,i} = \{B'_{t,i,1}, \dots, B'_{t,i,T}\}$ represents the expected consumption profile for consumer i given that a reduction request for time period t was issued. $B'_{i,j,t}$ represents the updated consumer demand at time t given that consumer j received a reduction request for time period i . This expected profile is generated using the consumer’s response profile p'_t (see equation 3).

The $(T + 1)$ th row in the matrix represents the case where no reduction request is issued to the consumer, i.e., $B'_{T+1,i} = B_i$.

We introduce a $(T + 1) \times |\mathcal{I}|$ matrix of binary decision variables, \mathbf{K} , where $K_{t,i} \in \{0, 1\}$. $K_{t,i} = 1$, indicates that consumer i will be issued a demand reduction request for time t , except for the case where $t = T + 1$, where a value of 1 indicates that no reduction request will be issued for consumer i on the day in question. Our solution finds the optimal values of \mathbf{K} that minimise the cost of the aggregate consumption profile to the supplier.

We define the updated deficit, $\mathbf{d}' = \{d'_1, \dots, d'_T\}$ and the updated surplus $\mathbf{s}' = \{s'_1, \dots, s'_T\}$ as the difference between the threshold Γ and the sum over all consumers and all time periods of the products of the decision variables K and the matrix of updated consumption profiles, \mathbf{B}' .

$$d'_t = \max(0, \sum_{i,j} K_{i,j} B'_{i,j,t} - \Gamma_t) \quad (5)$$

$$s'_t = \max(0, \Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t}) \quad (6)$$

We now find the values of K that minimise the cost to the supplier, as defined in the following objective function:

$$\mathbf{K}^* = \arg \min_{\mathbf{K}} \sum_t^T s'_t \rho_{sell}^t + d'_t \rho_{buy}^t \quad (7)$$

subject to the following constraints:

$$d'_t \geq 0 \quad (\text{constraint 1})$$

$$d'_t \geq \sum_{i,j} K_{i,j} B'_{i,j,t} - \Gamma_t \quad (\text{constraint 2})$$

$$s'_t \geq 0 \quad (\text{constraint 3})$$

$$s'_t \geq \Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t} \quad (\text{constraint 4})$$

$$s'_t \leq \Gamma_t x_t \quad (\text{constraint 5})$$

$$s'_t \leq \Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t} + M(1 - x_t) \quad (\text{constraint 6})$$

$$\sum_{t \in T} K_i = 1 \quad (\text{constraint 7})$$

In order to formulate the problem as a mixed integer linear program, we must remove the non-linear *max* functions that appear in the definitions of d'_t and s'_t (see equations 5 and 6).

The definition for d'_t requires the variable to take a value greater than or equal to 0. Constraint 1 states this explicitly. Similarly constraint 2 states that the value of s'_t must take a value greater than or equal to the second argument in its *max* function, $\sum_{i,j} K_{i,j} B'_{i,j,t} - \Gamma_t$. We assume that the value of ρ_{buy} will always be positive (occasions where suppliers will be paid by the grid operator to use excess generation are extremely rare), these constraints are sufficient to guarantee that the value of d'_t will take the greater value of 0 or $\sum_{i,j} K_{i,j} B'_{i,j,t} - \Gamma_t$. Were it to take a value greater than this, the objective function could be minimised further by reducing the value of d'_t .

We begin the process for s'_t in the same way, with constraints 3 and 4 requiring the variable to take values are both greater than or equal to 0 and greater than or equal to $\Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t}$. However, since the value of ρ_{sell} may take a negative value (the refund paid by the balancing mechanism to the supplier for their surplus was greater than the price the supplier originally paid for the electricity), we need to add further constraints to prevent the value of s'_t from inflating. This would happen whenever ρ_{sell} took a negative value because the objective function would be minimised when the value of $s'_t = \infty$. To prevent this inflation, we must also include constraints 5 and 6. These constraints introduce a new binary variable $\mathbf{x} = \{x_1 \dots x_t\}$, where $x_t \in \{0, 1\}$. Constraint 5 states that the value of s'_t must be less than or equal to the value of $\Gamma_t x_t$ where Γ_t represents the upper bound on the value of $\Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t}$ (when the total updated power demand = 0). Constraint 6 introduces a big-M coefficient (a very large number), where $-M$ represents the lower bound of $\Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t}$ (where

the updated power demand = ∞). The effect of these two constraints means that x_t will take the value 0 when $\Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t} < 0$ and 1 otherwise and the value of s'_t cannot take a value larger than the greater of 0 or $\Gamma_t - \sum_{i,j} K_{i,j} B'_{i,j,t}$.

Constraint 7 requires that the sum of any row in the matrix \mathbf{K} must be equal to 1 (a consumer can only receive a single reduction request per day or no request). We apply this constraint to limit the inconvenience caused to any one consumer.

In order to identify the subset of consumers who, when targeted, result in the lowest expected cost to the supplier, every combination of consumers must be evaluated. This means the size of the solution space grows exponentially in the number of consumers in the pool. However, as we will show in the next section, the algorithm can find optimal solutions for consumer pools that number in the tens of thousands.

In the following section we simulate, using real-world data, a community of residential electricity consumers engaged in a demand response programme to evaluate the performance of our solution.

5. EMPIRICAL EVALUATION

In this section we provide an empirical evaluation of our consumer selection algorithm. We first introduce the real-world consumption dataset that we use to evaluate our algorithm, before describing our experimental setup, including how we are simulating customer responses. We then provide results that shows our approach is capable of performing peak reduction on real world data and that we can achieve these reduction involving only a portion of the entire pool. Finally, we show that the algorithm is capable of scaling to communities of tens of thousands of consumers.

5.1 HES Dataset

The Household Electricity Use Study⁴ (HES) was carried out on behalf of the Department for Environment, Food and Rural Affairs (DEFRA), the Department of Energy and Climate Change (DECC) and the Energy Saving Trust between April 2010 and April 2011. The study monitored the household appliance usage of 251 residential dwellings in the UK. 26 of these dwellings were monitored over a period of one year at a resolution of 10 minutes, whilst the remaining 225 dwellings were monitored over a rolling one month period at a 2 minute resolution. We evaluate our model on the dataset produced from this study.

5.2 Experimental Setup

In what follows, we will provide details of our experimental setup, including our pre-processing of the HES dataset, and our approach to synthesising consumer response profiles.

5.2.1 Processing of consumption data

For all of the experiments detailed in this section, we have divided each day into 12 two hour time steps. Time periods

⁴<https://www.gov.uk/government/collections/household-electricity-survey>




Figure 2: An example aggregate demand profile. The shaded areas show the days adjacent to the day that is undergoing demand response.

of 2 hours reflect a suitable granularity when dealing with human consumers and have been used for this purpose before [15].

In order to generate a large amount of unique real-world data, we have placed a sliding window on each household’s profile and extracted the profiles for each weekday (including the 12 hours immediately before and after each weekday) as a separate consumer, i.e. for a household h we have extracted day 2 and day 3 and will consider them as separate consumers. By doing this we are able to perform experiments on sets of up to 10,000 unique days, where each day represents a consumer.⁵ By splitting up a household’s data in this way, we risk generating multiple consumers with similar underlying behaviour. However, since a consumer’s contribution to a demand response event is a combination of both the consumer’s consumption and response profile, consumers exhibiting similar consumption will still provide varied contributions to the demand response effort.

In order to account for the possibility of a consumer shifting consumption to an adjacent day, we extract consumption data in chunks of 48 hours. For this reason many of the consumption graphs have time axes starting at time step $t = -6$, corresponding to noon on day $d - 1$, and ending at $t = 17$, corresponding to noon on the day $d + 1$. As we are performing demand response for day d , we are only interested in reducing peaks in demand that fall in the interval $t = 0$ and $t = 11$, which correspond to the 12 time steps of day d . Figure 2 shows an example demand profile and highlights the area corresponding to day d .

5.3 Simulation of Customer Responses

In order to test our model’s reaction to customer responses, we generate a number of synthetic response profiles. In order

⁵A disadvantage of this approach is that we are not able to show how special events (such as major sporting events or days with cold temperatures) impact on the aggregate demand profile, but this is a minor concern.




Figure 3: The performance of the ‘Response Aware’ reduction strategy. Forward market threshold, Γ is denoted by the dashed grey line. Performance was averaged over 20 runs and the error bars represent the 95% confidence interval.

to generate the pre-emption probabilities, we select a value representing the most likely shift for the consumer by sampling from a normal distribution mean = -1 , $\sigma = 1$. This biases this maximum probability to shifts near the reduction request timestep, to represent the tendency of consumers to shift their consumption immediately before or immediately after a demand response event [11]. In Figure 1, this maximum pre-emption probability is shown at $\tau = -2$. The remaining pre-emption probabilities trail away from this value. A similar approach is taken to calculating the deferral probabilities, with the most likely shift being found by sampling from a normal distribution with mean = 1 , $\sigma = 1$.

There is currently insufficient data relating to consumer’s responses to propose a detailed data-driven model. However, despite its simplicity, this model remains consistent with existing models of consumer convenience and with observations made during demand response field trials [11, 7].

In order to focus on the impact of our consumer response model on system performance, in the experiments detailed in this section, the consumer’s willingness to shift was set to 1 for all consumers (the consumer pool always complies with reduction requests from the supplier). The results presented below therefore represent an upper bound on the algorithm’s performance.

5.4 Simulation of an Electricity Market

We simulate an energy market with a number of residential consumers. For the purposes of clarity we set uniform prices such that $\rho_{buy}^t = 6 \times \rho_{sell}^t \forall t \in T$. The value of ρ_{buy} is so much greater than ρ_{sell} because the value of ρ_{buy}^t represents the punitive price per kWh imposed by the balancing mechanism and ρ_{sell}^t represents the difference between the refund per kWh paid by the balancing mechanism and the original price paid for that electricity by the supplier. A more realis-

tic price scheme may see the prices vary in line with the level of demand – high demand leading to higher values of ρ_{buy} and ρ_{sell} , however, we wish to show that the peak reduction performance our algorithm achieves is not dependent on an externally determined price scheme. Similarly, for clarity, we define a flat threshold, Γ , where Γ_t is equal to 90% of the maximum demand on day d .

5.5 Reduction Strategies

To evaluate the performance of our targeting algorithm, we consider three reduction strategies:

- Response Aware (‘Aware’): The algorithm utilises the consumers’ response profiles when computing the optimal subset of consumers to target with demand reduction requests.
- Response Ignorant (‘Ignorant’): The algorithm does not take consumers’ response profiles into account when computing the optimal subset of consumers to target. The algorithm only takes into account the shiftable load the consumer has available at the timestep under evaluation. This strategy will demonstrate how our novel targeting algorithm works with current consumer modelling techniques.
- Peak Period Only (‘Peak Only’): All consumers receive a demand reduction request for the time period with the highest load. This represents a classical approach to the problem of demand response, where a supplier will issue a blanket request to all participants in a demand response scheme [10]. No account is taken of a consumer’s likely response and no consideration is given to the amount of shiftable load available to a consumer.

In what follows, we evaluate these three strategies in terms of peak reduction, supplier cost reduction and in terms of the percentage of the consumer pool who are recruited by the scheme.

5.6 Peak Reduction Performance

We measure the impact that our solution has in terms of peak reduction by considering the maximum power demand value on the day under reduction before and after reduction. We calculate a percentage peak reduction as follows:

$$P_{red} = \frac{\max_t(B_t)}{\max_t(\sum_{i,j} K_{i,j} B'_{i,j,t})} * 100 \quad (8)$$

Figure 3 shows the average performance of the response aware reduction strategy. It shows that the targeting algorithm consistently manages to reduce the peak in time slot 10. We show a comparison of peak reduction performance in Figure 4. This shows that the additional information provided by the consumer response profiles leads to increased peak reduction when compared to the ‘response ignorant’ and naïve ‘peak period only’ reduction. The increased performance of the ‘response aware’ strategy over the ‘response ignorant’ strategy is statistically significant at a 95% level of confidence. The ‘peak period only’ strategy also occasionally leads to negative peak reductions – that is, the peak

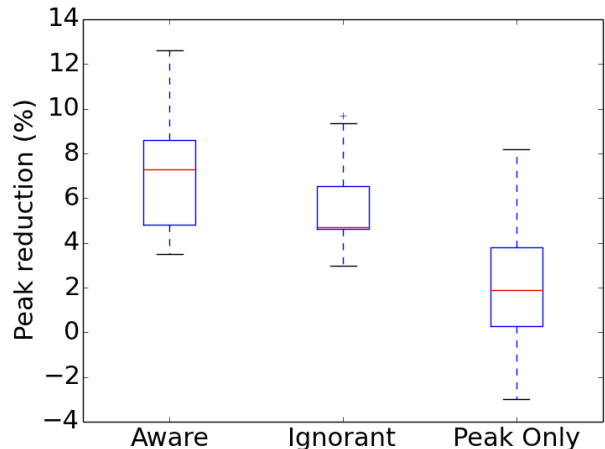


Figure 4: Peak reduction performance of the targeting algorithm for three different reduction strategies.

in demand grew after the demand response event. Figure 5 shows how the ‘peak period only’ reduction strategy has successfully reduced the original peak, but it has led to the creation of a secondary peak in the preceding timestep that is larger than the original.

Figure 6 shows that in addition to significant peak reduction, both the ‘response aware’ and ‘response ignorant’ strategies lead to substantial cost savings for the supplier. The difference in cost reduction between the ‘response aware’ and ‘response ignorant’ strategies is not statistically significant, but this result shows that the inclusion of consumer response behaviour does not negatively impact the potential cost savings produced by our targeting algorithm and so a supplier would have no economic reason for choosing to use the simpler reduction strategy.

5.7 Consumer Recruitment

One of the benefits of the approach we outline in this work, is that incentives will only be targeted to those consumers whose expected behaviour leads to a reduction in the cost to the supplier. On average, our ‘response aware’ solution issues demand reduction requests to 56.4% of the consumer pool compared to the ‘peak period only’ reduction strategy, which sends requests to 100% of the consumer pool. The ‘response ignorant’ reduction strategy recruited an average of 57% of consumers showing that our targeting algorithm is capable of reducing the recruitment of consumers even when the consumer behaviour model is omitted. This means that our targeting algorithm results in fewer interruptions to the daily lives of the supplier’s consumers, potentially minimising the onset of ‘demand response fatigue’, as well as leading to fewer incentives being paid to consumers whose behaviour following a reduction request actually leads to an increase in the supplier’s costs.

In Figure 7 we show how our solution distributes recommendations among the different time periods. Where a traditional approach to demand response, such as that repre-

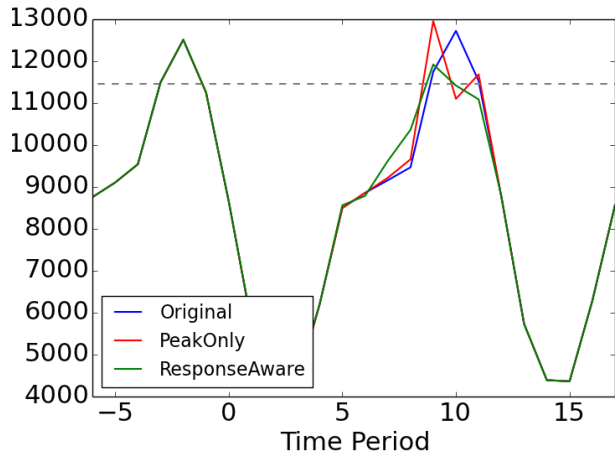


Figure 5: An example of the ‘payback effect’ generated by the ‘peak period only’ reduction strategy (red profile). The green profile represents the reduction performance of our ‘response aware’ solution on the same consumer pool. The forward market threshold, Γ is denoted by the dashed grey line

sented by the ‘peak only period’ strategy may only issue demand reduction requests to the period experiencing a peak in demand, our solution only issues slightly more than 50% of the recommendations to that peak period. This demonstrates that it is not only the selection of the right subset of consumers that contributes to the performance of the strategy, but the time periods that those requests specify also have a role to play.

5.8 Scalability

Finally, Figure 8 shows how the running time for the algorithm varies with an increasing number of consumers.⁶ These results show that although the size of the solution space increases exponentially with the number of consumers, solutions for large scale problems can still be found in reasonable time. We utilise IBM ILOG CPLEX Optimization Studio (CPLEX), which is an industry standard software package for mathematical optimisation, to solve the consumer targeting problem. The CPLEX pre-solver makes use of the branch and cut algorithm to identify opportunities to reduce the size of the problem. The branch and cut method utilises the branch and bound algorithm to solve a continuous relaxation of the linear program and makes use of cutting planes (additional linear constraints) to progressively tighten those relaxations. The addition of these ‘cuts’ normally reduces the number of branches requires to solve a MIP, enabling solutions to be found for larger scale problems.

The scalability of our algorithm means that it could potentially be used to coordinate all of a supplier’s consumers across a localised area such as a town/city, enabling the supplier and the grid operator to more efficiently manage demand across that area.

⁶The experiments were executed on an Intel Core i7-2600 workstation with 3.40GHz cores and 16GB of RAM.

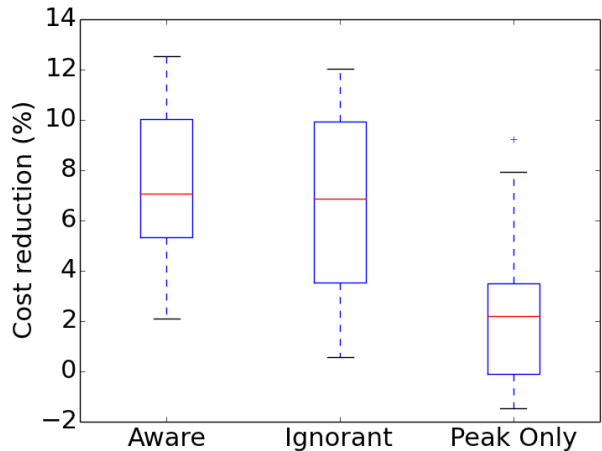


Figure 6: Cost reduction performance of the targeting algorithm for three different reduction strategies.

6. DISCUSSIONS AND FUTURE WORK

In this work we presented a rich model of customer behaviour that describes how consumers will respond to reduction requests. We also proposed a mixed integer programming solution to select the optimal subset of consumers that should be targeted with demand reduction requests. We evaluated this solution using electricity usage data from residential dwellings within the United Kingdom. We have shown that our solution is capable of performing significant peak reductions whilst reducing supplier incurred costs. We demonstrated that our consumer response model leads to improved peak reduction without compromising on the costs incurred by the supplier when compared to a consumer behaviour model that ignored consumer responses. We also evaluated our algorithm against a classical approach that saw all consumers within the consumer pool being targeted with a reduction request for the time period with the maximum demand usage and showed that our solution showed greater peak reduction and increased cost savings for the supplier. Finally, we have shown it to be capable of coordinating the electricity usage of tens of thousands of consumers with only a few minutes of computing time.

In future work we will propose a method to learn a consumer’s response profile using the system’s previous interactions with the consumer and evaluate how the performance of our algorithm changes when the consumer response profile contains uncertainty. We will also introduce uncertainty to the input, by utilising predictions of appliance usage, such as those provided by [15].

Acknowledgements

We thank the anonymous reviewers and Dr AJ Brush for their comments, which lead to improvements in this manuscript. This work was carried out as part of the ORCHID project (EPSRC ref: EP/1011587/1) and received support from the EPSRC-funded International Centre for Infrastructure Futures (EP/K012347/1).




Figure 7: The distribution of recommendations to different timeslots.




Figure 8: Running time of our targeting algorithm as the number of consumers varies. Error bars show a 95% confidence interval

7. REFERENCES

- [1] Assessment of customer response to real time pricing. Technical report, Edward J. Bloustein School of Planning and Public Policy, 2005.
- [2] M. Albadi and E. El-Saadany. A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989 – 1996, 2008.
- [3] V. Chandan, T. Ganu, T. K. Wijaya, M. Minou, G. Stamoulis, G. Thanos, and D. P. Seetharam. idr: Consumer and grid friendly demand response system. In *Proceedings of The 5th ACM International Conference on Future Energy Systems (e-Energy’ 14)*, number EPFL-CONF-198476, 2014.
- [4] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding. Progress in electrical energy storage system: A critical review. *Progress in Natural Science*, 19(3):291–312, 2009.
- [5] E. Costanza, J. E. Fischer, J. A. Colley, T. Rodden, S. D. Ramchurn, and N. R. Jennings. Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 813–822. ACM, 2014.
- [6] R. de Sá Ferreira, L. A. Barroso, P. Rochinha Lino, M. M. Carvalho, and P. Valenzuela. Time-of-use tariff design under uncertainty in price-elasticities of electricity demand: A stochastic optimization approach. *Smart Grid, IEEE Transactions on*, 4(4):2285–2295, 2013.
- [7] T. Ericson. Direct load control of residential water heaters. *Energy Policy*, 37(9):3502–3512, 2009.
- [8] J.-H. Kim and A. Shcherbakova. Common failures of demand response. *Energy*, 36(2):873–880, Feb. 2011.
- [9] J. K. Larsen. Correspondence regarding the rocky mountain power decision summary report on purpa time-based metering and communication standard. Technical report, 2007.
- [10] K. P. Ltd. Kiwi power: our service. Website. Accessed On 2015-01-12.
- [11] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 5–12. International Foundation for Autonomous Agents and Multiagent Systems, 2011.
- [12] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the ‘smarts’ into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [13] A. Rogers, S. Ramchurn, and N. R. Jennings. Delivering the smart grid: Challenges for autonomous agents and multi-agent systems research. In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012), Toronto, CA, July*, pages 22–26, 2012.
- [14] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox. Homeostatic utility control. *Power Apparatus and Systems, IEEE Transactions on*, (3):1151–1163, 1980.
- [15] N. C. Truong, J. McInerney, L. Tran-Thanh, E. Costanza, and S. D. Ramchurn. Forecasting multi-appliance usage for smart home energy management. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 2908–2914. AAAI Press, 2013.
- [16] E. UK. Wholesale electricity market report: Winter season: to end of january 2014. Technical report, 2014.
- [17] J. Xiao, J. Y. Chung, J. Li, R. Boutaba, and J.-K. Hong. Near optimal demand-side energy management under real-time demand-response pricing. In *Proceedings of Network and Service Management (CNSM), 2010 International Conference on*, pages 527–532, 2010.