

# Adaptive Incentive-based Demand Response with Distributed Non-Compliance Assessment

Gururaghav Raman<sup>a,b</sup>, Bo Zhao<sup>c</sup>, Jimmy Chih-Hsien Peng<sup>a,\*</sup>, Matthias Weidlich<sup>c</sup>

<sup>a</sup>*Department of Electrical and Computer Engineering, 2 Engineering Drive 3, National University of Singapore, Singapore, 117581*

<sup>b</sup>*Singapore-ETH Centre, Future Resilient Systems, CREATE campus, 1 CREATE Way, #06-01 CREATE Tower, Singapore, 138602*

<sup>c</sup>*Institut für Informatik, Humboldt-Universität zu Berlin, Berlin, Germany, 10099*

---

## Abstract

Traditional residential incentive-based demand response (DR) programs use fixed incentive structures that do not incorporate closed-loop feedback to compensate for non-compliance by participants. In practice, such programs may not reliably meet their event goals. To address this challenge, real-time feedback can be used to adaptively modify the participants' incentives, an approach which has not been proposed before. This paper proposes a flexible monitoring framework to detect potential non-compliance, whereby a second DR event is adaptively scheduled with higher incentives. In this context, constraints are presented to prevent over-compensation and gaming of the DR system by the participants. This novel dual-event design is implemented using a distributed event-stream monitoring framework to preserve scalability and ensure low monitoring costs. The merits of the proposed DR design are demonstrated at a utility-scale for 100,000 residents, while also considering the adoption of residential electric vehicles that are poised to increase the flexibility of the demand in the distribution system.

*Keywords:* demand response (DR), distribution system, residential energy consumption control

---

\*Corresponding author

*Email address:* [jpeng@nus.edu.sg](mailto:jpeng@nus.edu.sg) (Jimmy Chih-Hsien Peng)

---

## 1. Introduction

Demand response (DR) programs play a crucial role in increasing the reliability of power distribution systems by leveraging consumers' demand flexibility to manage the supply-demand gap [1, 2, 3]. Incentive-based, or event-based DR programs require changes in the consumers' demand patterns—usually, a load reduction during a specified event period—in exchange for monetary incentives [4]. Event details, including the event timing and incentive, are communicated to the DR participants before the actual event period.

However, real-time feedback about their performance during the event may not be available to the consumers or used by the utility [5]. Any data collected by the utility about the consumer performance is only used for post-event analyses, primarily for settlement (see Fig. 1(a)). In this open-loop implementation, the expected quantum of peak energy savings may not materialize due to multiple factors, e.g., hot weather resulting in higher cooling demand [6, 7].

While numerous studies have shown the importance of real-time consumer feedback in eliciting short and long term behavioral change in consumers [8, 9, 10], such feedback mechanisms have surprisingly not been incorporated widely by utilities. Some forms of feedback have been introduced in small field trials [4, 9], but the incentive offered to the participants was kept unchanged regardless of the observed performance during the DR event. Others have applied feedback in the context of optimal scheduling of heating or cooling loads enrolled in direct load control programs [11, 12, 13], or for price-based DR where automated flexible loads are controlled based on the dynamic price signal from a real-time market, e.g., in the transactive system proposed by the Pacific Northwest National Laboratory [14]. However, all of these designs explicitly require consumers to install home energy management systems and smart appliances capable of remote control, resulting in high costs and privacy concerns [6, 7] and leading to low adoption rates. Therefore, the above closed-loop DR designs are not applicable to residential consumers who respond manually to curtail their

consumption, which is the case in an incentive-based DR program.

The absence of DR feedback limits the utility in two ways. First, an inadequate performance cannot be detected *before* the event period ends. Second, the utility lacks flexibility in terms of how trade-offs—incentives to consumers vs. coping with an inadequate response and hence an inadequate peak demand reduction—are managed. In other words, since the incentive offered to the DR participants cannot be adaptively changed to improve their performance during the event, the utility is forced to rely on other resources to make up for any slack in the demand response. This is not economical and results in the under-utilization of the demand flexibility.

To address these challenges, this paper proposes monitoring of the DR performance in real-time to predict potential non-compliance of the system to the DR goals during the DR event. In conjunction, the proposed scheme increases the incentive offered to the participants to annul any deficit in the response. Different to existing studies and based on [15], the higher incentive is implemented using a *second* DR event, rather than the more impractical alternative of continuously varying the incentive in a closed loop. The second DR event happens within the time frame of the first event, as illustrated in Fig. 1(b). This design is simple, and as demonstrated later on, is effective in achieving the expected peak energy reduction. As such, the proposed DR design preserves the traditional schema of a DR event and can therefore be implemented only using existing DR infrastructure. Since the monitoring scheme should be scalable at the utility-level, this paper employs a distributive event-stream approach [16] wherein the monitoring system creates and responds to specific events of interest rather than processing continuously-generated real-time measurements. This further reduces the computational and communication costs.

In summary, the contributions of this paper are as follows:

1. It proposes a real-time monitoring system for traditionally open-loop incentive-based DR programs to robustly achieve the expected peak energy reduction.

2. It proposes an event-stream monitoring-based architecture for predicting potential non-compliance of the system to the DR goals before the DR period ends. This monitoring is adaptive and distributed, thereby reducing the communication and computational overheads.
3. It proposes a novel dual-DR-event design, wherein once non-compliance is predicted for the system, a second DR event is created to provide additional incentive to the participants in order to nullify the energy reduction deficit. Such a DR design has never been proposed previously in the literature. Constraints are also derived to prevent over-compensation and gaming.
4. A utility-scale demonstration of the proposed DR design is presented to illustrate the benefits of the proposed DR design while also considering the impact of increasing residential demand flexibility due to EV adoption.

The rest of the paper is organized as follows. Section 2 describes the proposed incentive-based DR scheme. Section 3 presents a case study that illustrates the merits of the proposed DR scheme, and Section 4 concludes the study.

## 2. Proposed Monitoring Scheme and DR Design

In a traditional incentive-based DR program, the utility informs the DR participants of an upcoming event while specifying the timing, energy reduction required, and incentive for compliance. During the event period, those consumers who commit to participating in the event defer their energy usage. Once the event ends, the performance of all the participants is assessed, and the utility compensates consumers who have fully achieved the required reduction (Fig. 1(a)). Formally, each participant  $j$  is found to have complied with the DR task if the following condition is satisfied:

$$\sum_{t=t_{start}}^{t_{end}} (B_{j,t} - P_{j,t}) \Delta t \geq \text{Specified energy reduction}, \quad (1)$$

where  $t_{start}$  and  $t_{end}$  respectively denote the starting and ending times of the DR event,  $B_{j,t}$  refers to the baseline load profile for that consumer (i.e., the fictitious

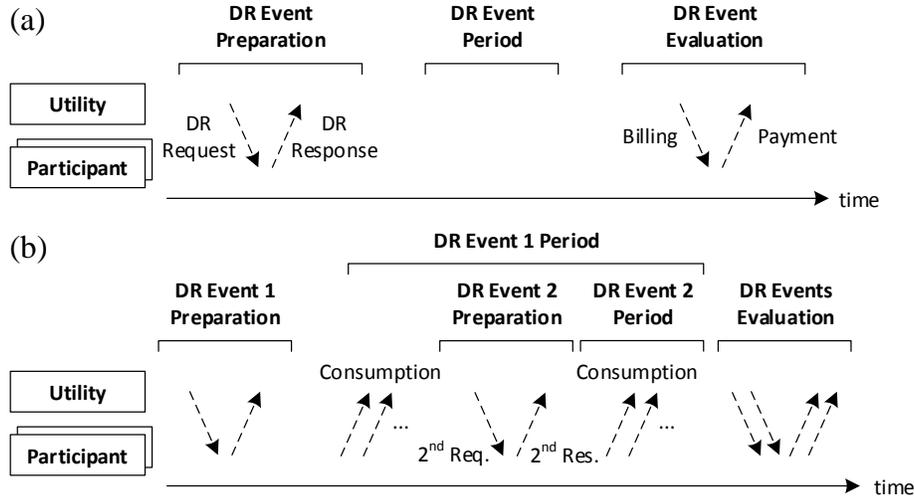


Figure 1: Incentive-based DR program design, showing: (a) the traditional single-event approach, and (b) the proposed dual-event design enabled by the real-time monitoring of participants' responses.

demand of the consumer had there been no DR event), while  $P_{j,t}$  is the measured load profile. The term  $\Delta t$  represents the metering resolution. Note that the above energy-based constraint is the 'Baseline Type-I' DR measurement and verification (M&V) technique defined by the North American Energy Standards Board (NAESB) Business Practice Standards [17] and is widely adopted by utilities in practice, e.g., see [4, 8, 5].

As mentioned previously, the utility traditionally does not monitor the performance of the DR participants in real-time during the event, and as a consequence, may not be able to assess the overall performance before the event ends. To address this, we propose to monitor the responses of DR participants during the event as shown in Fig. 1(b). Using this, the utility can predict the success of the DR event even as it progresses. If it is predicted that the required performance at the system-level can be achieved, no further action is taken, and all the participants who have successfully completed their individual DR reductions are compensated for their responses. However, if the utility predicts that the performance may fall short at the system-level, it has the flexibility

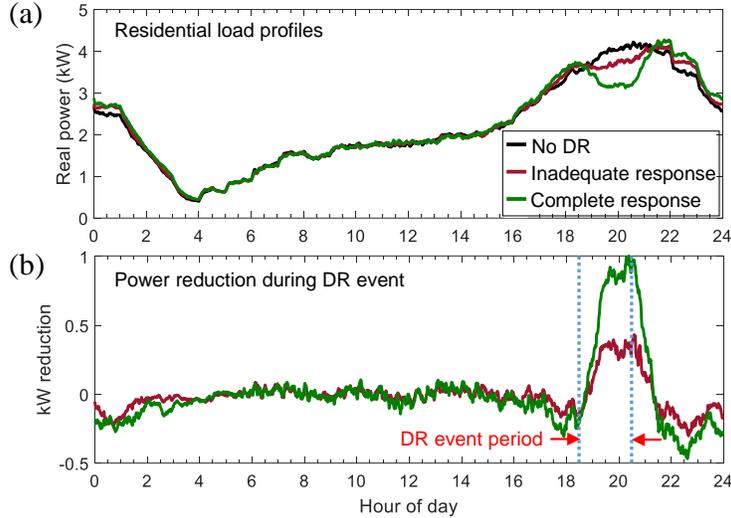


Figure 2: Inadequate vs. complete demand response of a residential consumer.

to modify the incentives of the participants in order to increase their response. The following subsection presents one such compliance prediction technique. We remark here that this specific technique is only an illustrative example and can be replaced in practice with any other method deemed suitable by the utility.

### 2.1. Real-time Monitoring and Compliance Assessment

Consider the average load profile of a typical residential consumer shown in Fig. 2(a), which was generated based on the probabilistic generation technique validated in [18] and [19]. The black plot in Fig. 2(a) corresponds to the case when no DR event exists. Consider the scenario where a DR event is scheduled from 6:30PM to 8:30PM; the green plot in Fig. 2(a) represents the consumer's typical response for a given incentive. Now, say for some reason (e.g., hot weather) that for a similar DR event on a different day, the consumer does not respond adequately, resulting in the load profile depicted in brown. The difference between the load profile with DR and that with no DR is plotted in Fig. 2(b). This figure indicates that unless the behavior of the consumer changes partly through the event period, the utility can predict if the consumer

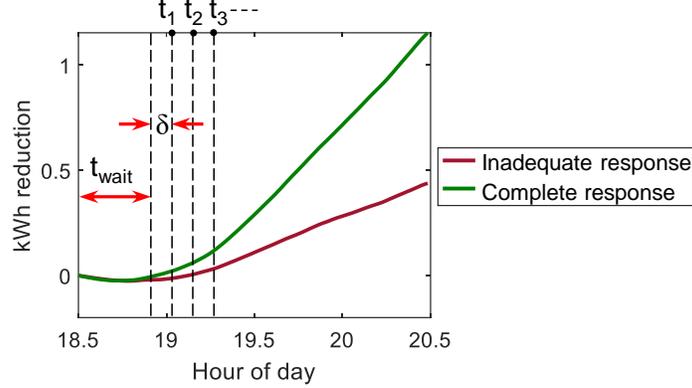


Figure 3: Profile of the cumulative energy reduction from the baseline during a DR event.

would be able to meet their event goal fully by monitoring the historically expected/complete (green) and measured (brown) profiles from the start of the event period. Based on this observation, we now derive a non-compliance prediction algorithm.

We begin by plotting the energy reduction profiles for the two cases, viz., inadequate response and complete response, for the event period. These are shown in Fig. 3. Note that these curves were derived by integrating the power reduction profiles shown in Fig. 2(b). These plots indicate that if the actual response (in terms of kWh reduction achieved) is lower than the expected/complete response at the beginning of the DR event, it is very likely that the energy curtailment at the end of the DR event would also be insufficient. This leads to the proposed non-compliance prediction algorithm, also based on the ‘Baseline Type-I’ M&V technique by the NAESB [17]:

1. Wait for a fixed time  $t_{wait}$  after the DR event begins for the pre-event rebound (see the undershoot in Fig. 3) to disappear.
2. Monitor the energy reduction at a constant interval  $\delta$  at times  $t_{test} = t_1, t_2, t_3, \dots$  such that  $t_{test} < t_{end}$  to verify if the following equation is satisfied for a participant  $j$ :

$$\sum_{t=t_{start}}^{t_{test}} (B_{j,t} - P_{j,t}) \Delta t < \sum_{t=t_{start}}^{t_{test}} (B_{j,t} - \hat{P}_{j,t}) \Delta t. \quad (2)$$

Note that  $[t_{start}, t_{end}]$  is the DR event period,  $B_{j,t}$  the baseline demand under no DR,  $P_{j,t}$  the measured demand during the event, and  $\hat{P}_{j,t}$  the load profile for the expected/complete response. This equation can be simplified into the following:

$$\sum_{t=t_{start}}^{t_{test}} P_{j,t} > \sum_{t=t_{start}}^{t_{test}} \hat{P}_{j,t} . \quad (3)$$

3. Predict non-compliance if eqn. (3) is satisfied for  $m$  consecutive testing times.

*Remark 1:* Transitioning from using eqn. (1) to (3) allows the prediction of non-compliance *before* the DR event ends.

*Remark 2:*  $B_{j,t}$  and  $\hat{P}_{j,t}$  are estimated using standard DR measurement and verification techniques based on historical performance data and/or randomized control trials [17].

*Remark 3:* The above algorithm waits for at least  $m$  consecutive testing intervals before issuing a prediction (see step 3 above). This is because in reality, the measured and expected load profiles would suffer from high variability as opposed to the average profiles presented in Fig. 2. Therefore, repeated testing is performed in order to avoid false predictions of non-compliance. Values for the parameters of this algorithm, i.e.,  $t_{wait}$ , the testing interval  $\delta$ , and  $m$  need to be tuned for a given system, e.g., using historical data from the DR program’s planning and design phase, before accurate predictions can be obtained.

To perform this monitoring process in a scalable, traceable, and privacy-preserving manner with low latency, we adopt an event-stream architecture [16]. For details on how the compliance-prediction algorithm can be implemented on smart metering streams, see [15]. In brief, a monitoring query operates on the stream of the consumers’ smart meter readings. Once non-compliance is predicted, a non-compliance event is added to an output stream of this query, which is then monitored by the utility in order to predict overall system-level non-compliance. For illustration purposes, we adopt the following simple rule for this: if more than  $\gamma_{ind}$  fraction of the total individual DR participants is

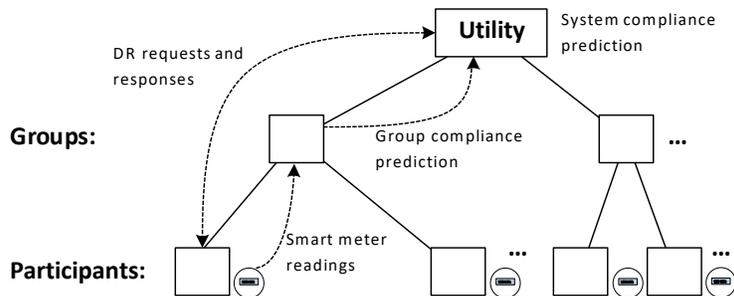


Figure 4: The proposed distributed DR monitoring system for groups of consumers. Individual consumers’ smart meter streams are monitored at the group level, and group-level compliance status is used by the utility for system-level compliance assessment.

predicted to be non-compliant, then the system-level performance is predicted to be non-compliant as well.

However, such individual-level testing and centralized monitoring may experience several challenges. First, it is difficult to accurately predict the expected reduction profile  $\hat{P}_{j,t}$  given the inherent variability of a single consumer’s demand. Second, the centralized nature of this monitoring would result in high communication and computation costs, while also being vulnerable to single-point failures. As such, a more accurate, efficient, and scalable approach is to assess the compliance for *groups of consumers* in a distributed fashion and use group-level compliance predictions to assess the system-level compliance as illustrated in Fig. 4. Importantly, this would also address an issue mentioned previously: while individual consumers may potentially change their performance during the event resulting in erroneous predictions of (non-)compliance, grouping multiple consumers’ demand in the analysis would mitigate the variability in their event performance.

A convenient and practical clustering approach is to group together all the consumers under a Load Aggregator. Once the groups are defined, the compliance of each group  $G_k$  is assessed similar to the procedure defined previously,

but instead of (3), we now test the following condition:

$$\sum_{j \in G_k} \sum_{t=t_{start}}^{t_{test}} P_{j,t} > \sum_{j \in G_k} \sum_{t=t_{start}}^{t_{test}} \hat{P}_{j,t} . \quad (4)$$

If (4) is satisfied for  $m$  consecutive testing intervals, the group’s response is predicted to be unsuccessful. Subsequently, system-level compliance is determined as follows: if more than  $\gamma_{group}$  fraction of groups is predicted to be non-compliant, then the overall system is predicted to be non-compliant.

In the proposed DR implementation, the monitoring rate at the various levels of the system is reduced once the compliance assessment is complete. For instance, consider the case where a group of DR participants is predicted to be non-compliant at an instance  $t_{nc}$  during the DR event. For this group, further compliance testing is not performed during the event period, and the monitoring rates for smart meters in this group are lowered to the rate required for settlement purposes. Therefore, even if the smart meter streams were initially monitored at a higher rate, say, at 1 min intervals, the rate can be reduced to 15 min intervals after non-compliance prediction. Meanwhile, the monitoring in the other groups remains unaffected, and the system-level monitoring query continues to predict the system-compliance status. Once the overall system is detected to be non-compliant, the monitoring rate for all the smart meters is then reduced. This approach significantly reduces the monitoring cost and allows scalable implementation in larger systems, as will be demonstrated in Section 3.4.3.

## 2.2. Utility Intervention during an Unsuccessful DR Event

Say the utility predicts an unsatisfactory response when consumers are given an incentive  $\lambda_o$ . It then offers a higher incentive  $\lambda_{new}$  for consumers to offset the deficit in the peak energy reduction. Here, rather than dynamically changing the incentive, we propose the incentive change to be implemented as a second DR event (see Fig. 1(b)). This way, by preserving the schema of the DR event, the proposed design allows for easier implementation by the utility. While the second event is announced at short notice, participants only need to reconsider

their energy usage in the immediate future, and therefore does not involve much preparation. Importantly, we do not implement more than two DR events to avoid the risk of participation fatigue [20], which may deter consumers from participating in future DR events.

### 2.2.1. Incentive for Second DR Event

Determining the new incentive  $\lambda_{new}$  requires the utility to have some knowledge about the behavior of the consumers in the system, which is a reasonable expectation [14]. One such technique is presented in the context of a case study in Section 3. In this subsection, constraints are presented so as to prevent over-compensation and potential gaming of the DR system by the participants.

First, while the second event task is offered to *all* participants, the proposed system uses a fixed probability  $p_{inc}$  to randomly select a *subset* of the participants that are compliant in the second event, who are actually paid the additional incentive. Indeed, such a lottery-based reward scheme has been found to elicit more response than a fixed reward [4], while also reducing the overall cost of the second event to the utility. The value of  $p_{inc}$ , or equivalently, the number of participants  $\kappa = (p_{inc} \times \text{total participants})$  actually offered the higher incentive is determined as follows:

$$\text{Hardware cost} + \text{Monitoring cost} + \kappa (\lambda_{new} - \lambda_o) \leq C_{saved}. \quad (5)$$

This equation compares the cost and benefit of the proposed scheme, with  $C_{saved}$  referring to the cost savings achieved by reducing the peak load through the second DR event.  $C_{saved}$  may also refer to the cost of using alternates such as direct load control to compensate the deficit in DR, which would allow the utility to compare the usefulness of the proposed approach vis-à-vis these alternatives. Here, the hardware cost is zero as the proposed scheme only involves making software changes (e.g., the implementation of new compliance monitoring queries and incentive-design method) to the utility's existing DR system and does not require any new hardware. More specifically, the utility already

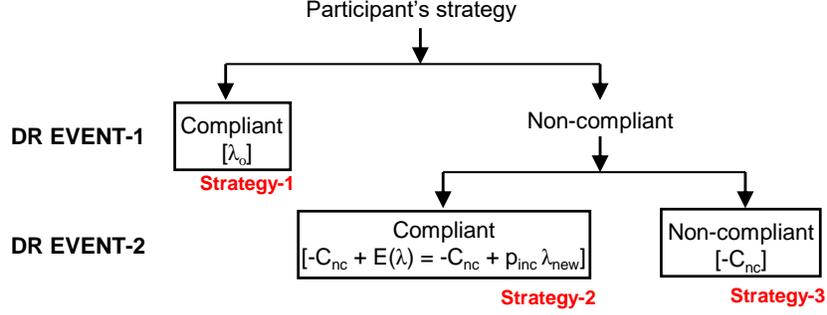


Figure 5: DR participants' expected payoffs (in square braces) under three possible strategies. Strategy-3 is strictly dominated, whereas Strategy-1 is the preferred strategy from the utility's perspective.

has in place smart meters and DR messaging infrastructure for communicating DR tasks to the consumers—the same systems can be utilized to schedule the second event as well. Therefore, the condition for profitability for the proposed DR design is obtained from (5) as:

$$\kappa \leq \frac{C_{saved} - \text{Monitoring cost}}{(\lambda_{new} - \lambda_o)}. \quad (6)$$

Second, to ensure participants provide a compliant response during the first event and do not wait for higher incentive during the second event (i.e., game the DR program), a penalty  $C_{nc}$  is imposed on participants who are non-compliant in the first event. Fig. 5 shows the expected payoffs under different possible scenarios. This figure shows that by stipulating:

$$\lambda_o > -C_{nc} + p_{inc} \lambda_{new}, \text{ or} \quad (7)$$

$$C_{nc} > p_{inc} \lambda_{new} - \lambda_o, \quad (8)$$

the utility can ensure that the dominant strategy for each participant is to fully participate in the initial event (i.e., Strategy-1) and thereby maximize their expected payoffs.

### 2.2.2. Mapping Additional Incentive to Consumer Response

The responses of the various participants to the second DR event would not be uniform. On the one hand, consumers who respond to the first event

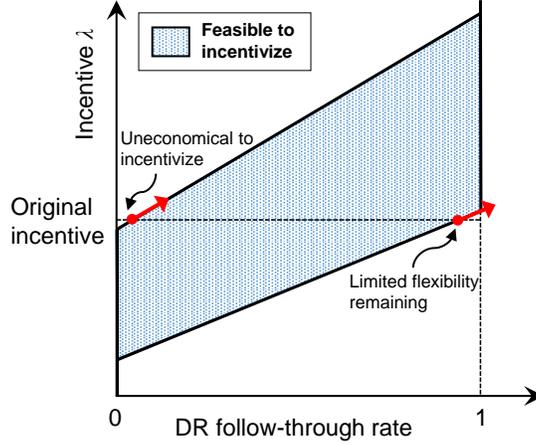


Figure 6: Mapping the incentive offered by the utility to the consumer response.

with high enthusiasm may not have additional flexibility left to provide, and therefore, any additional incentive may not result in meaningful returns. On the other hand, consumers who do not respond at all to the first event may be reluctant to respond during the second event as well, even for an increased incentive. These are respectively represented by the lower and upper bounds of the hatched region in Fig. 6, considering an affine mapping between the incentive offered and the response (similar models have been widely adopted, especially for thermostatic loads, e.g., see [14, 21]). The second event targets those consumers in the middle of the above spectrum, i.e., those whose characteristic lies within the hatched region in Fig. 6. As we show later in Section 3.2, a majority of consumers in reality likely fall into this category, underscoring the usefulness of the proposed DR design.

### 3. Performance Evaluation

This section presents simulation results evaluating the performance of the proposed dual-event DR design using a case study. The simulation procedure and case study are first described, followed by results highlighting the merits of the proposed scheme.

### 3.1. Simulating Residential Load Profiles

This subsection describes how the load profiles for the home appliances and the EVs in the system were obtained in our simulations.

#### 3.1.1. Home Appliance Load Profiles

The load profile of the home appliances for each of the residences is simulated in a bottom-up fashion, using the model and specifications introduced and validated in [18, 19]. In brief, this approach uses the probabilities of starting an appliance at a given time of the day to generate the load profile for the residence. The appliance-use timings are randomly generated at each time step based on the above probabilities, and the sum of the individual powers results in the residential load profile. If a DR event is requested by the utility, the appliance-use probabilities are modified to reflect the consumers' actions in deferring their energy usage away from the specified DR event period. Here, a variable  $\theta_{ft} \in [0, 1]$  is used to define the degree of follow-through of the resident, with  $\theta_{ft} = 1$  indicating no use of deferrable appliances (e.g., washers, dryers, and dishwashers) during the DR event period, and  $\theta_{ft} = 0$  meaning that the resident takes no action to change their energy usage behavior. Mathematically, say a deferrable appliance has a probability  $p_t$  of starting at a time  $t$ . Then, if the time step  $t$  falls within the DR event period, the revised starting probability of that appliance at time  $t$  is given by the following:

$$p_t^* = p_t (1 - \theta_{ft}). \quad (9)$$

This deferred probability,

$$\sum_{t \in \text{DR event period}} (p_t - p_t^*), \quad (10)$$

is then added to a random time outside the event period so as to model the deferred load being used at other times of the day.

#### 3.1.2. Residential EV Charging Load Profiles

As for the EV profiles, the simulation approach presented in [22] is adopted to generate load profiles when no DR event is requested by the utility. Based

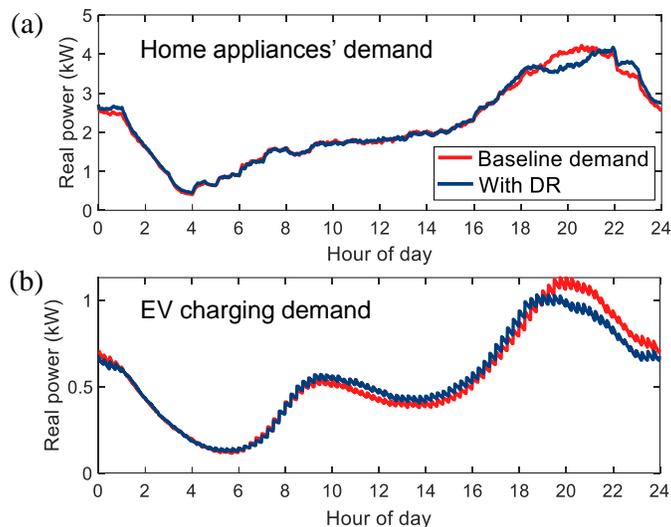


Figure 7: Average load profile for: (a) home appliances in one residence, and (b) charging one residential EV.

on one of the largest trials of residential EVs in Europe, this study proposes a bottom-up approach, where the behavior of each EV user (and thereby, the load profile of that EV) is parameterized by the following variables: the probability distribution of the number of charging events per day, probability that each charging event begins at a given time of the day, and lastly, the probability distributions of the initial and final state-of-charge values for the EV battery. When an EV user participates in a DR event, the user defers the EV charging away from the specified DR event period. In our simulations, this behavior is modeled as a reduction in the probability of a charging event commencing during the DR event period, similar to the home appliances' usage as described by eqns. (9) and (10).

Fig. 7 illustrates the resulting average load profiles of appliances in one residence and an EV while considering a DR event between 6:30-8:30PM, with  $\theta_{ft} = 0.3$  as an example.

Table 1: Simulation Parameters

DR event		Monitoring		Participant grouping	
Parameter	Value	Parameter	Value	Group	Consumers
$t_{start}$	6:30PM	No. of groups	5	G1	1-10,000
$t_{end}$	8:30PM	$t_{wait}$	20 min	G2	10,001-30,000
		$\delta$	5 min	G3	30,001-35,000
		$m$	3	G4	35,001-60,000
		$\gamma_{group}$	0.4	G5	60,001-100,000

### 3.2. Case Description

For the case study, we consider a system with 100,000 residential consumers, all of whom are assumed to be enrolled in the utility’s incentive-based DR program. Residents in the system are randomly assigned EVs depending on the EV penetration level. This is taken here as 22.48%, which is the 2030 forecast for EV adoption in the UK. As explained in Section 3.1, the response of each of these consumers is parameterized by a follow-through rate  $\theta_{ft} \in [0, 1]$ . In our study, to ensure that the simulations are realistic, values for the consumers’ follow-through rates are obtained from an online survey [23]. Specifically, the survey participants were shown a DR message, and asked to specify how likely they were to follow-through on such a message. Their responses indicated an average  $\theta_{ft} = 0.45$ , which aligns with metrics reported from previous DR field trials such as [4]. Importantly, about 52% of respondents indicate a follow-through rate of  $0.5 \leq \theta_{ft} < 1$ , suggesting that they fall within the hatched region in Fig. 6 and therefore can meaningfully contribute during the second event. In our simulations, the follow-through rate of each household is assigned to a randomly-chosen survey response.

Subsequently, the load profiles for each of the residences in the system are simulated. In particular, the behavior of the residents (including their appliance use and residential EV charging) is simulated in Matlab, and the event-stream monitoring setup is implemented using the Esper engine. The overall approach

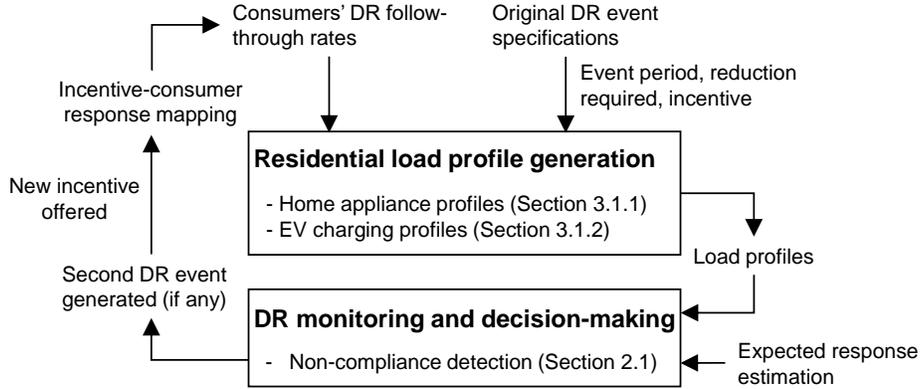


Figure 8: Flowchart illustrating how the proposed dual-event DR scheme was simulated.

is illustrated in Fig. 8, and the simulation parameters are listed in Table 1. Here, the parameters  $t_{wait}$ ,  $\delta$ ,  $m$ , and  $\gamma_{group}$  were selected empirically so as to accurately detect non-compliance for a 20% reduction in follow-through rates of consumers in this case study. We note here that a limitation of our survey was that we could not capture the participants' incentive-response mapping, which varies based on the location, demographics, and other factors. We therefore consider a set of different maps, M1-M4, which are depicted in Fig. 9(a) and study the performance of the proposed DR design for each. These mapping functions are such that for the original incentive offered by the utility (denoted by  $\lambda_o$ ), the resultant follow-through rate for each consumer is their expected follow-through rate  $\theta_o$  (i.e., when their response is complete), which in our simulations is derived from the survey.

### 3.3. Determining the New Incentive for the Second DR Event

Consider a scenario where on a particular day, the consumers' characteristics shift upwards as shown in Fig. 9(b), leading to a reduced follow-through rate  $\theta_{red}$  for the same incentive  $\lambda_o$ . The task for the utility now is to determine what new incentive  $\lambda_{new}$  would mitigate this deficit in the response. To this end, the utility can leverage the system-wide cumulative energy reduction profiles (similar to those depicted in Fig. 3) to approximate the mean reduced follow-

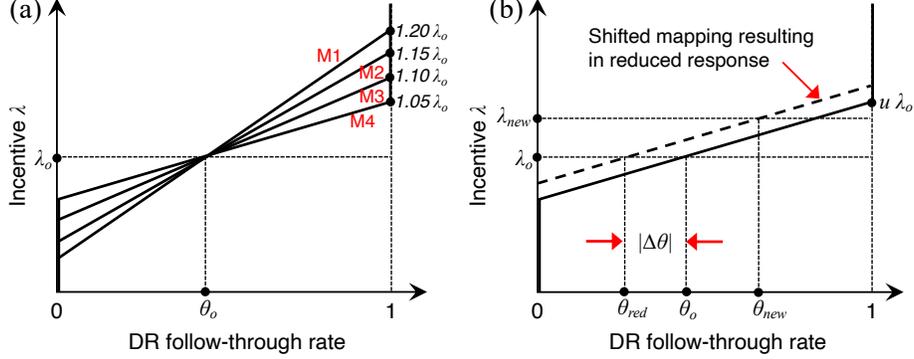


Figure 9: (a) Incentive-consumer response mappings considered in the case study. (b) Reduction in the consumer response.

through rate of all consumers,  $\bar{\theta}_{red}$ :

$$\frac{\bar{\theta}_{red}}{\bar{\theta}_o} \approx \frac{\text{Measured energy reduction } (t_{nc})}{\text{Expected energy reduction } (t_{nc})}, \quad (11)$$

where  $t_{nc}$  is the time at which system non-compliance is predicted, and  $\bar{\theta}_o$  is the mean expected follow-through rate (i.e., equal to 0.45, which is the mean follow-through reported in our survey). Here, the assumption is that the ratio of the measured and expected energy reduction profiles is approximately equal to that of the mean reduced and expected follow-through rates. Subsequently, the utility can construct a linear mapping similar to that shown in Fig. 9(b) between the incentive offered and the mean system-wide follow-through rate  $\bar{\theta}_{red}$  to obtain the new incentive:

$$\lambda_{new} = \lambda_o + \frac{(u-1)\lambda_o}{1-\bar{\theta}_o} (\beta \bar{\theta}_o - \bar{\theta}_{red}). \quad (12)$$

Here,  $u > 1$  is defined such that  $u\lambda_o$  is the point where the linear characteristic under normal conditions intersects with the line representing unity follow-through rate (see Fig. 9(b)). The factor  $\beta > 1$  is used as an *overcompensation* during the second event so as to mitigate the initial deficit in the response (i.e., because  $\bar{\theta}_{red} < \bar{\theta}_o$ ). In effect,  $\beta \bar{\theta}_o$  becomes the targeted follow-through rate during the second event so that during the overall DR event, the expected energy reduction can be achieved.

To simulate how individual consumers react to the new incentive  $\lambda_{new}$ , we use the characteristic presented in Fig. 9(b) to determine the new follow-through rate for each consumer:

$$\theta_{new} = \theta_{red} + \frac{(\lambda_{new} - \lambda_o)(1 - \theta_o)}{(u - 1)\lambda_o}. \quad (13)$$

By substituting (12) in (13), we derive the following:

$$\theta_{new} = \theta_{red} + \left( \frac{1 - \theta_o}{1 - \bar{\theta}_o} \right) (\beta \bar{\theta}_o - \bar{\theta}_{red}). \quad (14)$$

This shows that by selecting (12) as the new incentive, the final follow-through rates of consumers are independent of the initial incentive  $\lambda_o$  and incentive-response parameter  $u$ , indicating the robustness of the proposed incentive mechanism.

### 3.4. Performance Evaluation

#### 3.4.1. Effectiveness in Achieving Peak Energy Reduction

Consider a scenario where unexpectedly, the participants' follow-through rates are reduced by 40%. As a result, with no changes to the incentive, the peak energy reduction reduces from the expected 83.28 MWh to 47.89 MWh; see Fig. 10 for the corresponding demand profiles. Here, the proposed group-based monitoring scheme analyzes the real-time demand and predicts that only two of the five groups are expected to be compliant at 7:11PM. Given the threshold  $\gamma_{group} = 0.4$ , it then predicts that the overall system would be non-compliant, and schedules a DR event beginning ten minutes later with a new increased incentive calculated according to (12) with  $\lambda_o = 1$  and  $\beta = 1.8$ . The resultant demand profile is then determined based on the new follow-through rates from (13) and is illustrated in Fig. 10. We observe that the energy reduction achieved by the proposed technique is 82.6 MWh, which is within 1% of the expected value of 83.28 MWh, thereby demonstrating the effectiveness of the proposed DR design. Recall that incentive design in (12) results in follow-through rates that are independent of the parameter  $u$  and by extension the specific incentive-response mapping function, which is why separate results are

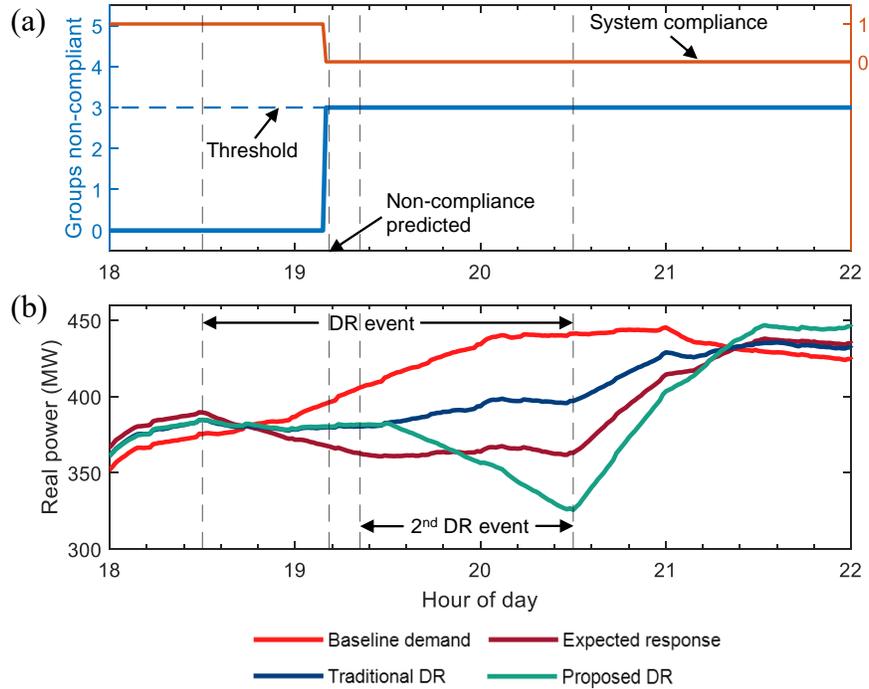


Figure 10: Illustrating the proposed DR scheme for a DR event from 6:30PM-8:30PM. (a) System non-compliance predicted at 7:11PM. (b) System demand for the baseline scenario with no DR, expected DR, traditional single-event response (reduced), and the proposed dual-event DR design.

not shown for M1-M4. We note here that the above results pertain to one simulation instance where participants' follow-through rates and load profiles were generated randomly as detailed in Section 3.1. Here, note that while individual consumers' demand profiles have variability across different simulation instances, when aggregated at the group (each comprising 5,000–40,000 consumers) or system level (100,000 consumers), the variability in the net demand profiles is small. Therefore, for computational simplicity, all results hereon are also presented for a single simulation instance.

Next, we simulate the system for varying levels of reduction in the follow-through rates and compare the energy reduction achieved by the proposed DR design with that of the traditional single-event design. Referring to Fig. 11, the

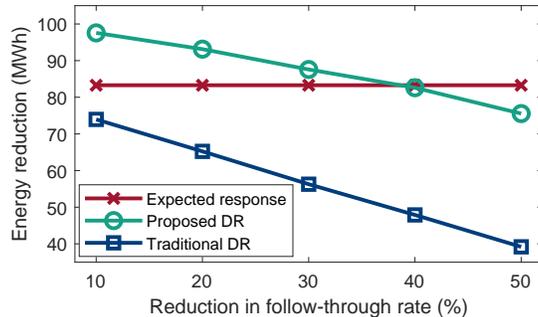


Figure 11: Comparing the effectiveness of the proposed dual-event DR design with the traditional single-event response.

proposed incentive determination technique may under- or over-compensate the peak energy reduction deficit. Further, there is a limit to the compensation it can provide—for instance, for a 50% reduction in the follow-through rates, there still remains a 9.3% deficit in the response. This can be partly ascribed to the simplistic method used in this study to determine the new incentive  $\lambda_{new}$ ; in practice, a utility could use more advanced methods such as adaptively tuning the factor  $\beta$  in order to achieve perfect compensation. Regardless, there does exist a practical limit to the amount of response that can be elicited during the second event, as we address later in Section 3.4.2.

*Speed of Non-Compliance Detection:* The effectiveness of the proposed DR design is dependent on the speed of non-compliance detection. This is illustrated in Fig. 12 while assuming the same event parameters as in Fig. 10. Notably, the faster the detection of non-compliance, the larger the window to mitigate the deficit in response. Conversely, a higher incentive needs to be provided to consumers to achieve the same performance when non-compliance is detected late; these results are not presented here due to length constraints.

*Impact of Increasing Demand Flexibility:* The net residential demand flexibility in the system is strongly dependent on the uptake of EVs by the residents [22]. To study how varying EV adoption impacts the benefits offered by the proposed DR scheme, we refer to Fig. 13. Systems with higher EV adop-

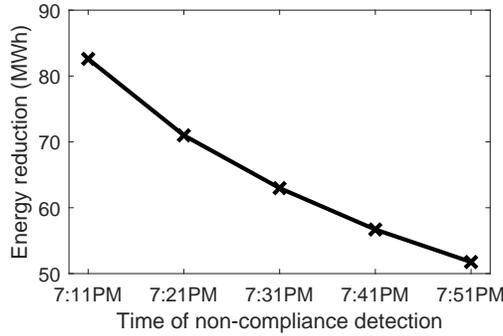


Figure 12: Variation in the DR performance with the speed of non-compliance detection. All other parameters except the time of non-compliance detection remain the same.

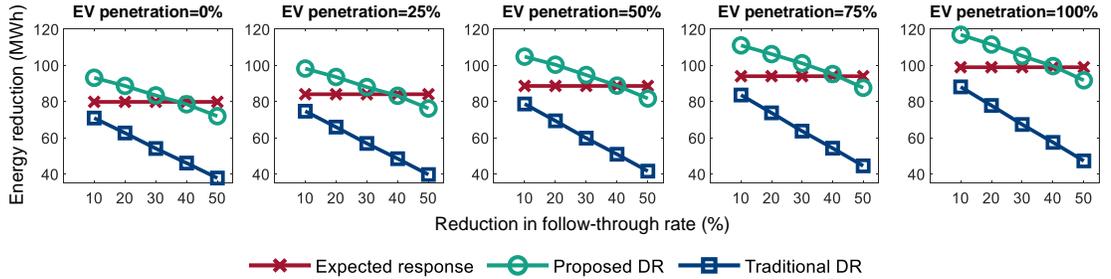


Figure 13: A comparison of the performance of the proposed and traditional DR designs as the demand flexibility increases.

tion have a greater need to avoid unsatisfactory DR, as the evidenced by the increasing difference between the energy reduction achieved by the proposed and traditional DR designs. For instance, for an EV penetration of 0%, the increased peak energy savings from the proposed DR design over the traditional DR design varies from 31-90%, while for 50% EV penetration, the improvement varies from 33-94%.

### 3.4.2. Economic Feasibility of Dual-Event DR

As shown in eqn. (6), there is a maximum limit on the total extra incentive ( $\kappa (\lambda_{new} - \lambda_o)$ ) paid during the second DR event for it to be economically feasible. Here, we determine exactly how much extra incentive would fully compensate the energy reduction deficit for our case study while considering

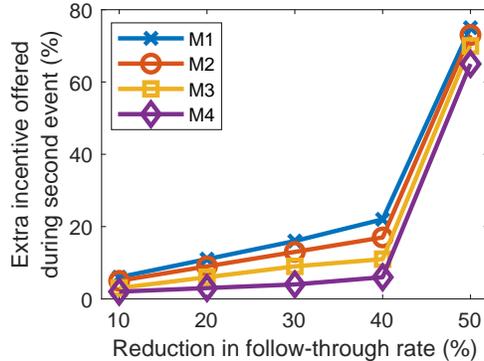


Figure 14: Additional incentive required during the second DR event to nullify the energy reduction deficit.

the four different incentive-consumer response mappings presented in Fig. 9(a). These results are shown in Fig. 14. We observe that a higher incentive is required when consumers are less engaged (mapping M1) when compared with those who are more responsive (mapping M4). Furthermore, we find that the rate at which the additional incentive increases grows dramatically when the initial consumer response reduces further and further. For instance, depending on the mappings M1-M4, the extra incentive ranges from 4-16%, 6-22%, and 65-75% for 30%, 40%, and 50% reduction in the follow-through rates, respectively. Beyond a point in the x-axis in Fig. 14, it becomes impossible to achieve the promised reduction for any incentive offered, i.e., even if all consumers were to have unity follow-through rate during the second event, it is not possible to meet the energy reduction goal. In such cases, the proposed incentive-based response would need to be supplemented with alternatives such as direct load control.

### 3.4.3. Scalability

To illustrate the scalability of the proposed DR scheme, we present the communication and computation costs involved in the monitoring process. Specifically, the communication cost is defined as the number of event transmissions, while the computation cost is defined as the number of events (i.e., queries executed over the various event streams) processed by the monitoring system.

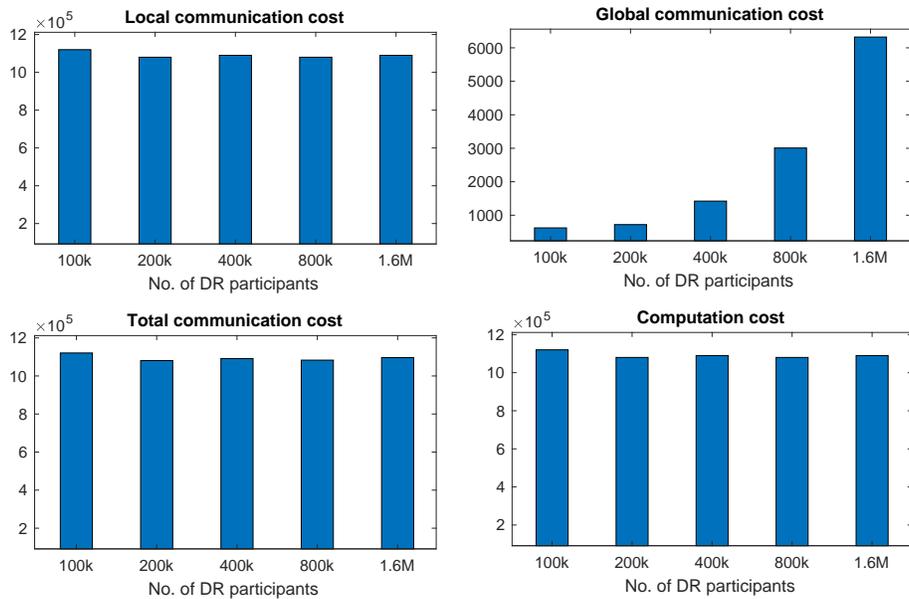


Figure 15: Communication and computation costs as the number of DR participants increases.

We simulate systems where the number of DR participants is increased from 100,000 to 1.6 million. The results are presented in Fig. 15, assuming every 10,000 consumers to form one group. Evidently, the monitoring costs remain nearly constant as the system size increases. This is because in the proposed distributed monitoring system depicted in Fig. 4, the monitoring of the participants is local within the group, allowing for parallel compliance assessments within the various groups. Global communications are only required between the group-level monitoring system and the utility. Therefore, to scale out the system is to actually add more groups, which only increases the number of global communications by the number of newly-added groups. Furthermore, the resolution at which the smart meter streams are monitored are reduced from 1 min intervals to 15 min intervals after non-compliance detection, which this further minimizes the communication costs.

#### 4. Conclusion

An effective DR program is only fair for all stakeholders, be it the utility or the consumers. To this end, this paper has proposed a flexible monitoring framework for assessing the performance of the DR participants in real-time, and in the case of potential non-compliance to the event goal, adaptively creating an additional event with higher incentives for participation. This is implemented using a distributed event-stream monitoring approach, which allows for scalable realization with low computation and communication overheads. It does not require new hardware and can be integrated easily into the DR management system already in place, and can be modified by utilities to suit their respective customer bases. Constraints have been derived to avoid the potential gaming of the DR system by participants, and to determine the profitability of the proposed scheme vis-à-vis alternatives such as direct load control. The merits of the proposed DR design were demonstrated for a system of 100,000 residential consumers, using bottom-up simulations of home appliances and electric vehicles, and realistic consumer behavior models.

#### Acknowledgements

The research was conducted at the Future Resilient Systems at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation Singapore. This research is supported by the National Research Foundation Singapore (NRF) under its Campus for Research Excellence and Technological Enterprise (CREATE) programme.

#### References

- [1] E. T. Hale, L. A. Bird, R. Padmanabhan, C. M. Volpi, Potential roles for demand response in high-growth electric systems with increasing shares of renewable generation, Tech. rep., National Renewable Energy Laboratory (2018).

- [2] H. C. Gils, Economic potential for future demand response in germany— modeling approach and case study, *Applied Energy* 162 (2016) 401–415.
- [3] Q. Wang, C. Zhang, Y. Ding, G. Xydis, J. Wang, J. Østergaard, Review of real-time electricity markets for integrating distributed energy resources and demand response, *Applied Energy* 138 (2015) 695–706.
- [4] M. Jain, et al., Methodologies for effective demand response messaging, in: *IEEE Int. Conf. Smart Grid Commun.*, 2015, pp. 453–458.
- [5] Oracle, Delivering on the smart grid promise: How Exelon utilities create value for every customer with dynamic pricing, <https://www.oracle.com/customers/exelon-bge-1-opower.html> (2016).
- [6] S. Z. Attari, M. L. DeKay, C. I. Davidson, W. B. De Bruin, Public perceptions of energy consumption and savings, *Proc. National Academy of Sciences* 107 (37) (2010) 16054–16059.
- [7] L. V. White, N. D. Sintov, Inaccurate consumer perceptions of monetary savings in a demand-side response programme predict programme acceptance, *Nature Energy* 3 (12) (2018) 1101.
- [8] V. Tiefenbeck, et al., Real-time feedback promotes energy conservation in the absence of volunteer selection bias and monetary incentives, *Nature Energy* 4 (1) (2019) 35.
- [9] R. Smale, G. Spaargaren, B. van Vliet, Householders co-managing energy systems: space for collaboration?, *Building Research & Inform.* 47 (5) (2019) 585–597.
- [10] N. Good, Using behavioural economic theory in modelling of demand response, *Applied Energy* 239 (2019) 107–116.
- [11] Y.-J. Kim, L. K. Norford, J. L. Kirtley, Modeling and analysis of a variable speed heat pump for frequency regulation through direct load control, *IEEE Trans. Power Syst.* 30 (1) (2014) 397–408.

- [12] K. McKenna, A. Keane, Residential load modeling of price-based demand response for network impact studies, *IEEE Trans. Smart Grid* 7 (5) (2016) 2285–2294.
- [13] M. Ye, G. Hu, Game design and analysis for price-based demand response: An aggregate game approach, *IEEE Trans. Cybern.* 47 (3) (2016) 720–730.
- [14] Pacific Northwest National Laboratory, Transactive system, <https://tinyurl.com/y3d5xxcf> (Dec. 2017).
- [15] G. Raman, J. C.-H. Peng, B. Zhao, M. Weidlich, Dynamic decision making for demand response through adaptive event stream monitoring, in: *Proc. IEEE Power Energy Soc. Gen. Meeting, 2019*, pp. 1–5.
- [16] M. Dayarathna, S. Perera, Recent advancements in event processing, *ACM Computing Surveys (CSUR)* 51 (2) (2018) 33.
- [17] M. L. Goldberg, G. K. Agnew, Measurement and verification for demand response, *Tech. rep.*, US Department of Energy USA (2013).
- [18] L. Chuan, A. Ukil, Modeling and validation of electrical load profiling in residential buildings in Singapore, *IEEE Trans. Power Syst.* 30 (5) (2015) 2800–2809.
- [19] G. Raman, J. C.-H. Peng, T. Rahwan, Manipulating residents’ behavior to attack the urban power distribution system, *IEEE Trans. Ind. Informat.* 15 (10) (2019) 5575–5587.
- [20] P. Ben-Nun, Respondent fatigue, *Encyclopedia of survey research methods* 2 (2008) 742–743.
- [21] K. Subbarao, et al., Transactive control and coordination of distributed assets for ancillary services, *Tech. rep.*, Pacific Northwest National Laboratory (2013).

- [22] J. Quirós-Tortós, L. Ochoa, T. Butler, How electric vehicles and the grid work together: Lessons learned from one of the largest EV trials in the world, *IEEE Power Energy Mag.* 16 (6) (2018) 64–76.
- [23] G. Raman, B. AlShebli, M. Waniek, T. Rahwan, J. C.-H. Peng, How weaponizing disinformation can bring down a city’s power grid, *PLOS ONE* 15 (8) (2020) e0236517.