Price Perturbations for Privacy Preserving Demand Response with Distribution Network Awareness

Constance Crozier, Aisling Pigott, and Kyri Baker

***Abstract*—Demand response (DR), where electricity consumption is shifted in response to incentive signals, can ease the transition to renewable generation. However, when many devices simultaneously respond to these signals there is the potential for violating local network constraints. Many of the proposed solutions require consumers to hand over control of their devices to third-parties. Here, we propose a method of using price distortions to coordinate distributed resources, which is both robust to local constraints and privacy preserving. We formulate the price distortion setting problem as a mixed-integer linear programming problem. Conditions are derived under which the method can guarantee constraint violations are eliminated within one time step. The method was tested using case studies involving both electric vehicles and smart heating/cooling systems. We show that the proposed method, under a scenario with maximum DR participation, can achieve 98% of the theoretical lower bound on the number of constraint violations. Furthermore, our method out-performs the benchmark where some devices opt-out of DR.**

NOMENCLATURE

Parameters

κ	Penalty of network constraint violation (\$/kW)
λ	Transmission-level real-time price (\$/kWh)
Δ_t	The size of one time step (hrs)
η	The efficiency of a device (%)
Θ	The safe power limit of the network (kW)
τ	The deadline by which a load must be met
C	Thermal capacitance of building (J/K)
d	The total inflexible demand at a bus (kW)
E	The energy requirement of a device (kWh)
I	Global horizontal irradiance (W/m^2)
k	Effective window area of building (m^2)
M	A large number which exceeds other parameters
N_t	The number of steps considered in the horizon
p	The power a device draws when on (kW)
R	Thermal resistance of building (W/K)
T_{max}	The maximum acceptable temperature (°C)
T_{min}	The minimum acceptable temperature (°C)
T^{out}	The temperature outside (°C)

Variables

σ	Violation of a distribution network constraint
ξ	Distribution price adjustment (\$)
c	The price below which a device will turn on (\$)
T	The internal temperature of a building (°C)
x	Binary variable indicating device on/off status

C. Crozier is with the School of Industrial & Systems Engineering, Georgia Institute of Technology. A. Pigott, and K. Baker are with the Dept of Civil, Environmental & Architectural Engineering, University of Colorado Boulder.

Sets

$\mathcal{J}_i$	The set of devices connected to bus i
$\mathcal{J}_T$	The set of thermal devices
$\mathcal{J}_D$	The set of deadline driven devices
$\mathcal{J}_U$	The set of uninterruptible devices
$\mathcal{S}_{xfm}$	The set of transformer thermal constraints
$\mathcal{S}_v$	The set of voltage constraints
$\mathcal{S}_b$	The set of branch thermal constraints

Notation

$^{(i)}$	At distribution bus i
$^{(j)}$	From smart device j or its associated building
$^{(s)}$	Relating to distribution constraint s
$\hat{}$	Estimate of the quantity
$_t$	At time step t

I. INTRODUCTION

DEMAND response (DR) broadly refers to altering power demand to achieve economic or grid-level goals. This is distinct from historic operation of the power grid, which has focused on altering power supply to match demand. Renewable generation presents increasing challenges for transmission system operators (TSOs), as they need to cope with greater uncertainty on the supply-side. The proliferation of devices such as electric vehicles (EVs) and smart thermostats in distribution networks has resulted in a large amount of latent flexibility being present at the low voltage level [1]. This presents an opportunity to improve the operation of the power system, without incurring additional infrastructure cost [2].

Dynamic pricing schemes are in early stage deployment, where dynamic tariffs (which represent the TSO system-wide cost) are passed down to consumers (e.g. [3]). These schemes incentivize consumers to align their demand with times of surplus renewable generation. However, real-world trials have shown that sending homogeneous prices to large number of flexible devices can increase local peak demand – creating problems for the distribution network [4]. Shifting demand to improve the TSO operation, while coordinating devices to prevent distribution system constraint violations is often termed the TSO-DSO coordination problem [5].

There has been some academic appetite for dynamic electricity prices which also reflect local distribution system operation (e.g. [6]). However, there are ethical and technical concerns with such an approach, and thus far no utility has successfully deployed spatially varying prices [7]. Instead, most previous works addressing this issue seek to include distribution system operation in the optimization constraints.

Some works have re-framed the economic dispatch problem to consider co-optimization of the transmission and distribution problems; e.g. including distributed solar generation [8], EV chargers [9], or management of battery storage systems [10]. However, these assume a centralized control structure which is unlikely to scale to millions of devices. A common method for reducing the complexity of the problem, is to use third-parties (often called aggregators) to interface with the TSO (e.g. [11]). Aggregators can co-ordinate devices and interface directly with the TSO, however additional controls may be necessary to ensure that distribution constraints are protected (e.g [12]). Additionally, passing control of their devices to a third-party aggregator requires transmission of sensitive information, which may threaten consumer privacy [13]. Although encryption methods can be used to protect consumer privacy (e.g. [14], [15]), these incurs additional costs. Furthermore, in order to anonymize individual devices, the exact position of the device in the network must be occluded. This means that robustness to some constraints can not be guaranteed (such as local voltage and branch loading).

Other works have investigated decentralized methods for DR, for example using the alternating direction method of multipliers [16], gradient projection [17], or the water-filling algorithm [18]. These decentralized formulations mean that device requirements are kept private, however they still require forecast consumption to be broadcast and can take many iterations to converge. Others proposed methods utilize cryptography to preserve privacy; for example, using differential privacy [19] or aggregation methods [20] such that patterns in the dataset are shared without disclosing individual consumer data. In [21] the authors propose a method where the utility receives a recommendation of consumers to elicit demand response from, that guarantees differential privacy. It is unlikely privacy would be maintained if the locations were also provided. In [22] a privacy-preserving consensus algorithm is developed where nodes send projected states (protecting the privacy of their initial position) to neighbours. However, this involves bi-lateral communication so can be slow to converge.

An alternative approach is for the controller to use inference from previous behavior to estimate the behaviour of devices in response to price signals. In [23] reinforcement learning is used to learn the optimal incentive strategy based on consumers previous response to price signals. This paper did not investigate the inclusion of distribution network constraints, however reinforcement learning has been shown in other cases to be effective at improving distribution system operation [24]. The limitation of using a model-free approach, such as reinforcement learning, is that there can be no robustness guarantee – we have no assurance that the constraints will be met where possible. In [25] the authors propose a strategy for the control of A/C units which does not require the operator to know the outputs of the individual units. Instead the units provide only velocity signals (signifying their rate of change of output in response to price), which gives a proxy for the distance of the unit from its comfort bound. Although effective, this approach only works with a single device-type.

In this paper, we propose a distributed method for DR which requires devices to only share current consumption – visualized in Fig. 1 (c). In this method, a local operator applies perturbations to the TSO price signals with the objective of avoiding local network constraint violations. This avoids both the local constraint violations of the top down method (a) and the privacy issues with the aggregator model in (b). Furthermore, our solution adopts a minimum hardware approach by removing additional market players and communication infrastructure. This is highly desirable from the consumers' perspective, because additional players and hardware come at a financial cost, which is ultimately passed on to the consumer [26]. Our proposed method allows diversity in consumer preference, and remains effective even when many consumers opt-out of flexibility provision.

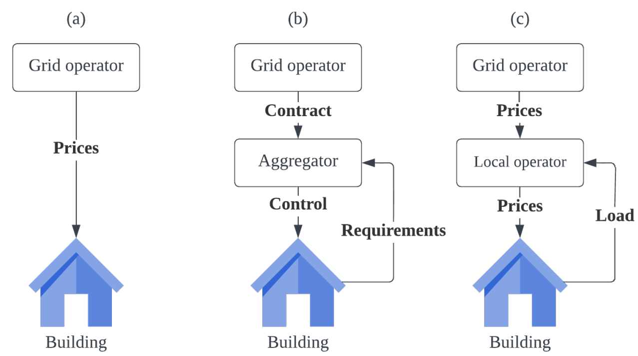

Fig. 1: Three architectures for exploiting distribution level resources for transmission level goals. (a) top-down: where TSO level price signals are passed straight to smart devices, (b) direct control: where a local aggregator takes control of the smart devices, (c) proposed method: where a local operator adjusts the TSO price signals in response to local observations.

The contributions of the paper can be summarized as follows. First, that we propose a method of DR that is distribution network aware and does not require iterative convergence. Second, that our method is privacy-preserving, meaning that users are not required to transmit their requirements or proposed energy consumption. Third, that the method does not require a 100% participation rate, meaning it is still effective even when some devices opt out of DR. Finally, that we release our simulation platform as open access – allowing other researchers to repeat our results and test similar methods.

II. Problem Formulation

A. Control Architecture

The proposed control scheme involves altering the transmission level price signals so as to avoid synchronization of local devices. We assume that the distribution system operator (DSO) defines these perturbations to the transmission system price. This concept is illustrated in an example network in Fig. 2. The locational marginal price (LMP), λ, is passed down from the transmission network and a number of smart devices j are connected via nodes i. Although we describe λ as an LMP in this manuscript, this approach would generalize to any flexibility service with a transmission level price, e.g. ancillary services. While DSOs have historically been predominately focused on network repairs, participation in real-time operation is becoming more common (e.g [3]).

The DSO defines the perturbations ξ (which can be positive and negative) that are used to diversify the price signal sent to each node. We consider two possible options in terms of implementation. The first option is that the consumer will pay the real time price signal $\lambda + \xi$, such that the perturbations influence the actual cost of their electricity. Previous works have suggested that such distribution system varying price signals can increase the economic efficiency of electricity supply [27] – particularly with large loads like EVs [28]. However, there are equity concerns that arise when sending local consumers different price signals, and in practice no utility has successfully implemented such as scheme [7]. The second option is that the price perturbations do not influence the price that the consumer pays – they still pay λ for their energy use. The DSO will design the price adders to minimize the cost the consumers pay using λ, hence saving consumers money and maximizing the exploited flexibility. This may be more practically implementable for regulatory reasons, but it may be challenging to incentivize consumer participation.

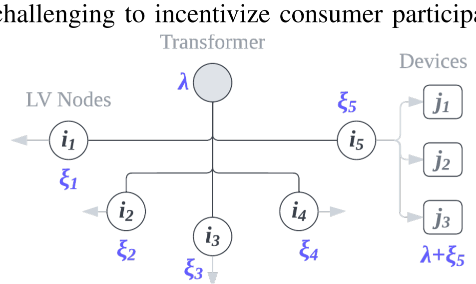

Fig. 2: An example network. Multiple low voltage nodes are connected to a transformer. All devices receive the price λ plus an additional offset ξ which is specific to their node.

In terms of communication, we assume that the DSO broadcasts the nodal price adders to devices, and that devices only share their current consumption. Importantly, the device requirements and intended consumption profiles are not shared with the DSO. Current consumption need not be communicated directly through the device, but rather through submetering at the smart meter level. An interesting further direction would be investigating whether the aggregate household consumption would be sufficient. It is worth noting that there are methods of load coordination which do not require this level of information. For example, using aggregated consumer data [29], however this approach can not include awareness of the distribution network. Alternatively, some use cryptography methods to preserve differential privacy (e.g. [15])) – but these methods still require device communication. Figure 3 shows an example communication between a newly activated device and the DSO controller. Iterations are not required, so this method represents a significantly reduced computational burden compared to existing approaches (e.g. [30]). Given that the DSO considers devices on a single circuit, scalability should not be a concern, but this will be explored further.

B. Devices

Here we consider a smart device j to be a single-speed device whose demand is delayable. Devices are defined by their power rating $p^{(j)}$ and efficiency $\eta^{(j)}$. We use the binary variable to $x_t^{(j)}$ to denote whether the device is drawing power

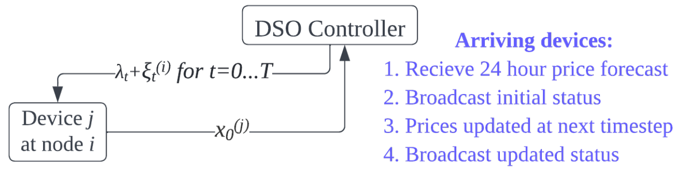

Fig. 3: The communication process when a device is activated.

at time t (a value of one implies that it is). Our formulation works for any smart devices which have communication abilities and can delay load. Many flexible loads already have this functionality, although it is worth noting that for smaller devices (e.g. charging of mobile phones) this functionality may be cost-prohibitive. We distinguish several categories of device, which have different constraints.

1) Thermal devices: These are electrical temperature controls – e.g. air conditioning units, heat pumps, and resistive heaters. These have inherent flexibility because consumers typically have a range of building temperatures which they find acceptable, and buildings have natural thermal storage. In this case the flexibility of the device is determined by its temperature bounds $T_{min}^{(j)}, T_{max}^{(j)}$; consumers can opt out of flexibility by setting an extremely narrow bound for acceptable temperature. Note that shifting thermal load will increase the total required energy, due to the thermal leakage [31].

2) Deadline devices: These are devices whose total energy consumption $E^{(j)}$ does not change when load is shifted forwards or backwards. Instead the flexibility of these devices are governed by a deadline $\tau^{(j)}$, the time by which they need to have been completed. A common example of a deadline device would be an EV charger.

3) Uninterruptible devices: These are a subset of deadline devices for which, once the device has turned on, the process must complete without interruption. This means that the flexibility is only in the start time – once the device has begun drawing power it has no more flexibility. An example is a dishwasher or washing machine with a delayed start.

C. User Interface

Previous research has indicated that simple choices for consumers elicits the optimal participation in DR programs [32]. Therefore, we have chosen the simplest possible user interface and input parameters. For thermal devices, consumers set the temperature bounds they are comfortable with. For deadline oriented devices, consumers set their deadline (the device will automatically communicate energy consumption).

In terms of DR scheme consumers have two or three options. For deadline-driven devices the consumer can choose: (1) to consume economically before a set deadline, (2) to consume immediately, or (3) to consume below a certain price. The second option allows consumers to opt out of DR, e.g. if their deadline is uncertain. The final option is for when the consumer does not necessarily need the energy by a deadline, but will consume power if the price falls extremely low (or negative). For thermal devices, there will be two options: (1) consuming economically, and (2) default consumption, which will minimize the total energy consumption of the device.

III. Control algorithms

Our proposed method involves two separate controllers – the device-level control which alters the device consumption in response to the incentive signal, and the DSO-level control which designs the price perturbations. We also define a benchmark case where an aggregator directly controls devices.

A. Model Predictive Controller

For devices j where the user has selected the economy setting, a local controller will need to optimize the device consumption. Here we assume that these devices will use model predictive control (MPC) to minimize their energy cost given the forecast set of prices. We define the forecast cost of device j's energy consumption as:

$$f^{(j)}(x) = p^{(j)}\Delta_t \sum_t (\hat{\lambda}_t + \hat{\xi}_t^{(i)})\ x_t^{(j)}\ , \tag{1}$$

where $x_t^{(j)}$ is a binary variable determining whether device j is consuming power at time t, $p^{(j)}$ is the device's grid-side power rating, Δ_t is the time step size, $\hat{\lambda}$ are the forecast transmission prices, and $\hat{\xi}$ are the forecast price adders. For deadline driven devices, the power profile must satisfy the energy constraint:

$$p\ \eta^{(j)}\ \Delta_t \sum_{t=0}^{\tau^{(j)}} x_t^{(j)} \geq E^{(j)}\ , \tag{2}$$

where $\eta^{(j)}$ is the efficiency of device j (given that p is defined from the grid-side), $\tau^{(j)}$ is the device's deadline, and $E^{(j)}$ is the total energy requirement. We use a greater than (rather than an equality constraint) because x are binary variables, so there is likely not a feasible solution to the equality constraint. Note that if bi-directional power flow is considered (such as with vehicle-to-grid chargers) then we need to define separate binary variables for charging and discharging and include an additional constraint on state-of-charge. For uninterruptible devices, we include the additional constraint:

$$x_t^{(j)} \geq x_{t-1} - \frac{1}{E^{(j)}} p^{(j)} \Delta_t \eta^{(j)} \sum_{m=0}^{t-1} x_m\ , \tag{3}$$

where the far right term gives the percentage of the required energy that has been met. Given that x are binary variables, this enforces that if $x_t = 1$ then subsequent x_t must be equal to one until the energy demand has been met. Thermal devices have different constraints, as they are governed by the buildings thermal properties. Here we use a linear RC model to determine to determine the building temperature change that results from electrical heating or cooling. These models are commonly used for this application (e.g. [33]). Specifically, we take our model and parameters from [34]. Therefore, we impose variables T for building temperature, which are linked in time by the following equation:

$$\begin{aligned} T_{t+1}^{(j)} - T_t^{(j)} = \Big(& \frac{T_t^{out} - T_t^{(j)}}{R^{(j)} C^{(j)}} + \frac{k^{(j)}}{C^{(j)}} I_t \\ & - \frac{\eta^{(j)}}{C^{(j)}} x_t^{(j)} p^{(j)} 10^3 \Big) 3600 \Delta_t\ , \end{aligned} \tag{4}$$

where $R^{(j)}$ is building thermal resistance, $C^{(j)}$ is the building thermal capacitance, T^{out} is the outdoor temperature, $k^{(j)}$ is the effective window area, and I is the global horizontal irradiance. The factors of 10^3 and 3600 are used to convert kW to W and /s to /hr respectively. Note that here $\eta^{(j)}$ will capture the conversion from electric to thermal power as well as any losses. In this case, given the short time horizon considered, we assume the outdoor temperature forecast to be known. The flexibility of the device is then given by:

$$T_{min}^{(j)} \leq T_t^{(j)} \leq T_{max}^{(j)}\ , \tag{5}$$

where $T_{min}^{(j)}, T_{max}^{(j)}$ are the minimum and maximum allowable building temperature for device j.

B. DSO controller

In contrast to the devices, the DSO controller is able to affect the prices, but has only estimates of device requirements. Additionally, we assume that the controller has no knowledge of future activating devices, so the controller must adapt to new devices as they are activated. This is particularly challenging for deadline-driven devices, which can add large loads unexpectedly. We assume that the controller knows the inherent properties of the device, such as its type, efficiency, and power rating. Given that the DSO can not dictate the device actions, we need a way of incorporating the price-sensitivity of loads. In order to maintain a linear formulation, we first define a new variable $c^{(j)}$, which is price signal below which the controller expects device j will consume power.

Remark 1. For deadline-driven devices j there will be a price signal $c^{(j)}$ below which the device will be on and above which the device will be off.

To see this, first consider that if we constrain the sum of a set of binary variables to be n and minimize a linear combination of them, the lowest cost n variables will be 1 and the remainder will be 0. This is analogous to devices optimizing their on/off behavior in response to a price signal. For devices selecting the 'consume immediately' option or uninterruptible devices which have already started $c^{(j)} \to \infty$. For devices selecting the 'consume below a certain price' option there will be a fixed value of c which does not vary with time. For other devices the value of $c^{(j)}$ will be constantly updated, according to the urgency of the device's constraints. However, the controller can gain some insight into price sensitivity by observing choices in the previous time step. Thermal devices are more complicated, given that the total energy consumed will increase if consumption is bought forward.

Remark 2. For thermal devices j there will be a price $\alpha_t c^{(j)}$ below which the device will be on and above which off.

Here we have introduced a factor α_t which will inflate earlier costs to account for the additional energy consumption caused by bringing forward heating or cooling. We assume:

$$\alpha_t = 1 + \frac{T_t^{out} - \hat{T}_{max}^{(j)}}{R^{(j)} \eta^{(j)}} 10^{-3} \Delta_t (N_t - t) \tag{6}$$

for cooling devices, and $(\hat{T}_{min}^{(j)} - T_t^{out})$ as the numerator for heating devices. This approximation is derived in the

Appendix. Therefore, considering a time horizon of N_t time steps of Δ_t, we formulate our DSO objective as:

$$f(.) = \sum_t^{N_t} \Big(\lambda_t \sum_j \hat{x}_t^{(j)} p^{(j)} + \sum_s \kappa^{(s)} \sigma_t^{(s)} \Big) + \sum_j c^{(j)} \quad (7)$$

where λ_t represents the transmission system price signal at time t, $\hat{x}_t^{(j)}$ is an estimate of j's charging profile. Therefore, this first term represents the estimated device energy cost using the transmission level pricing. The variable $\sigma_t^{(s)}$ is a slack variable describing the violation of the distribution constraint s at time t and $\kappa^{(s)}$ is a parameter describing the cost of violating that constraint. Therefore, this second term represents the penalty for violating the distribution system constraints. It is assumed that $\kappa \gg \lambda$, such that the optimization will prioritize minimizing constraint violations, and then maximizing the TSO response. Lastly, $c^{(j)}$ represents the maximum cost at which j will consume power, enforcing that consumers will individually act to minimize their cost. This final term does not need to be weighted, because there is no direct link between this and the other two terms, so the there is no trade-off.

The control variables to be chosen are the DSO price adders $\xi_t^{(i)}$, which dictate the addition to the transmission price which devices at bus i receives at time t – such that devices at that bus will receive signal $\lambda_t+\xi_t^{(i)}$. However, the formulation also includes decision variables for: estimates of device behaviors $\hat{x}$, estimates of the threshold costs c, and estimates for the building temperatures $\hat{T}$.

The constraints of the problem are as follows:

$$\sum_t \xi_t^{(i)} = 0 \quad \forall i \quad (8a)$$

$$(\lambda_t + \xi_t^{(j)}) \leq \alpha_t^{(j)} c^{(j)} + M(1-\hat{x}_t^{(j)}) \ \forall j,t \quad (8b)$$

$$\alpha_t^{(j)} c^{(j)} \leq (\lambda_t + \xi_t^{(j)}) + M\hat{x}_t^{(j)} \ \forall j, t < \hat{\tau}^{(j)} \quad (8c)$$

$$\alpha_t^{(j)} - 1 = 0 \quad \forall t, j \in \mathcal{J}_D \quad (8d)$$

$$\alpha_t^{(j)} - 1 = \frac{\hat{T}_{min}^{(j)} - T_t^{out}}{R^{(j)}\eta^{(j)}} 10^{-3}\tau_t \quad \forall t, j \in \mathcal{J}_T \quad (8e)$$

$$\hat{x}_0^{(j)} = x_0^{(j)} \quad \forall j \quad (8f)$$

$$\begin{aligned} \hat{T}_{t+1}^{(j)} - \hat{T}_t^{(j)} = \Big(&\frac{T_t^{out} - \hat{T}_t^{(j)}}{R^{(j)}C^{(j)}} + \frac{k^{(j)}}{C^{(j)}} I_t \\ &- \frac{\eta^{(j)}}{C^{(j)}} \hat{x}_t^{(j)} p^{(j)} 10^3 \Big) 3600\Delta_t \end{aligned} \quad (8g)$$

$$\hat{T}_{min}^{(j)} \leq \hat{T}_t^{(j)} \leq \hat{T}_{max}^{(j)} \quad \forall t, j \in \mathcal{J}_T \quad (8h)$$

$$\sum_t \hat{x}_t^{(j)} p^{(j)} \Delta_t \geq \frac{1}{\eta^{(j)}} \hat{E}^{(j)} \quad \forall j \in \mathcal{J}_D \quad (8i)$$

$$\hat{x}_t^{(j)} = 0 \quad \forall j \in \mathcal{J}_D, t > \hat{\tau}^{(j)} \quad (8j)$$

$$\hat{x}_{t+1}^{(j)} - \hat{x}_t^{(j)} \geq -\sum_{m=0}^{t} \frac{\hat{x}_m p^{(j)} \Delta_t \eta^{(j)}}{\hat{E}^{(j)}} \forall t, j \in \mathcal{J}_U \quad (8k)$$

$$\sigma_t^{(s)} \geq 0 \quad \forall s, t \quad (8l)$$

where we use $\mathcal{J}_D, \mathcal{J}_T, \mathcal{J}_U$ to denote the set of devices which are deadline driven, thermal, and uninterruptible respectively. The set $\mathcal{J}_i$ defines the devices which are connected to bus i. Constraint (8a) enforces that the price adders on bus i over the time horizon must sum to zero. This achieves fairness between devices on the network, because the adders must have zero bias at all locations – this prevents some areas on the network from receiving greater DSO intervention than others. Constraint (8b) and (8c) enforce the definition of $c^{(j)}$ through the big M formulation [35] – where M is some large number much greater than $\alpha_t c$, these constraints are equivalent to:

$$\begin{aligned} &x_t^{(j)} = 1 \quad \text{if } \lambda_t + \xi_t^{(i)} \alpha_t^{(j)} \leq c^{(j)} \\ &x_t^{(j)} = 0 \quad \text{otherwise}\,. \end{aligned} \quad (9)$$

Constraints (8d) and (8e) enforce the definition of the inflation factor $\alpha_t^{(j)}$ – ensuring that it is unity for deadline-driven devices, and the loss-compensated factor previously discussed for thermal devices. Here we have used the variable $\tau_t = \Delta_t(N_t - t)$ to denote the amount of time that will pass between time step t and the end of the horizon. Note that (8e) assumes a heating device, if the device is cooling $\hat{T}_{min}$ should be replaced with $\hat{T}_{max}$. Constraint (8f) includes the observation of the current state of the device. This provides the controller with some information about $c^{(j)}$; if the device is currently consuming power we know that $c^{(j)} \leq \lambda_0 + \xi_0^{(i)}$. Constraints (8g) and (8h) enforce the behavior of thermal devices; where the controller is working with estimates of the building temperatures and preferences. Constraint (8i) enforces the behavior of deadline-driven devices, although again with estimates of the energy requirements. The estimated deadline of the device is enforced by (8j), where the consumption of the device is set to zero after the estimated device deadline. Finally, constraint (8k) enforces the behavior of uninterruptible devices.

The constraints of the distribution system are captured through slack variables $\sigma_t^{(s)}$ which are defined as follows:

$$\sum_j \hat{x}_t^{(j)} p^{(j)} \leq \Theta + \sigma_t^s - \sum_i d_t^{(i)} \quad \forall t, s \in \mathcal{S}_{xfm} \quad (10a)$$

$$-\overline{b}^{(i)} - \sigma_t^s \leq \mathbf{F}_{[i,:]}(\hat{\mathbf{x}}p + \mathbf{d}) \leq \overline{b}^{(i)} + \sigma_t^s \quad \forall t, s \in \mathcal{S}_b \quad (10b)$$

$$\underline{v} - \sigma_t^s \leq \mathbf{M}_{[i,:]}(\hat{\mathbf{x}}p + \mathbf{d}) \leq \overline{v} + \sigma_t^s \quad \forall t, s \in \mathcal{S}_v \quad (10c)$$

where we used the notation $\mathcal{S}_{xfm}, \mathcal{S}_{br}, \mathcal{S}_v$ to define the of distribution constraints s which are related to transformers, branch limits, and voltage bounds respectively. Transformer constraints are enforced by (10a), where the thermal limit of the network is defined by a maximum power rating Θ. This could be easily extended to multiple transformers if the network considered has multiple voltage levels – the sum over i just needs to only include nodes under the transformer in question. Branch constraints are enforced by (10b), where the matrix $\mathbf{F}$ uses a loss-less approximation to map injections to branch flows, as described in [36]. Voltage constraints are enforced using (10c) where a linear approximation of the power flow constraints has been used to maintain a MILP formulation. The matrix $\mathbf{M}$ maps load injections to bus voltages. Although multiple such approximations exist, we suggest using [37], given that it is applicable to unbalanced three-phase distribution networks.

C. Direct Control Benchmark

In order to compare the action of the the proposed method to direct control from aggregators, we define the following benchmark method. We assume that a controller has direct control over the devices which have chosen to participate in DR (but no visibility or control over those who have opted-out). Here we use a modified version of the objective (7):

$$f(x) = \Delta_t \sum_{t}^{N_t} \left(\lambda_t \sum_{j} x_t^{(j)} p^{(j)} + \sum_{s} \kappa^{(s)} \sigma_t^{(s)} \right), \quad (11)$$

where the controller can directly choose x, so we have replaced $\hat{x}$ with x and removed the c term. Rather than the constraints defined in the previous section, we can apply the constraints (2)–(5) from the device MPC controllers, given that the controller knows the device constraints. Plus the distribution constraints (11) with the true rather than estimated values of x. We maintain the soft distribution constraint because for consistency we assume that the aggregator prioritizes meeting consumer demand over the distribution constraints.

IV. Constraint Robustness Guarantee

In the proposed formulation all distribution constraints are linear with device consumption. Therefore we can say that we can express any violation of constraint s at time t as:

$$\sigma_t^{(s)} = C + k^{(1)} x_t^{(1)} + k^{(1)} x_t^{(2)} \cdots + k^{(N_j)} x_t^{(N_j)}, \quad (12)$$

where C is a constant and $k^{(j)}$ is a continuous variable describing the impact of $x_t^{(1)}$ on $\sigma_t^{(s)}$. The value of k will vary based on the constraint. For example, for transformer constraints $k^{(j)} = p^{(j)}$, the rated power of the device. Whereas for branch and voltage constraints the value of k will depend on the properties of the network. Given that the controller can not forecast future arriving devices and has only estimates for device parameters, it is not possible to place bounds on the values of $\sigma_t^{(s)}$. However, we can place bounds on the controller response to a constraint violation. In other words, given some strictly positive $\sigma_0^{(s)}$, can we place bounds on the controller response $\sigma_0^{(s)} - \sigma_1^{(s)}$. From (12) it follows:

$$\sigma_0^s - \sigma_1^s = \sum_j k^{(j)} (x_0^j - x_1^j), \quad (13)$$

The expression $(x_0^j - x_1^j)$ can take one of three values: $-1, 0, 1$. Therefore, if all devices have flexibility and opt into DR the maximum reduction is given by:

$$\sigma_0^s - \sigma_1^s = \sum_{j \in \mathcal{J}_{on}^+} k^{(j)} - \sum_{j \in \mathcal{J}_{off}^-} k^{(j)} \quad (14)$$

where $\mathcal{J}_{on}^+$ is the set of devices with $k^{(j)} > 0$ which are on at $t = 0$, and $\mathcal{J}_{off}^-$ is the set of devices with $k^{(j)} < 0$ which are off at $t = 0$. The benchmark method, which has direct control of x, will achieve this reduction if the device constraints are not binding. Recall that the proposed method controls devices indirectly with price adders $\xi_t^{(i)}$. Intuitively we can see that the proposed method can achieve the same outcome as the benchmark if it can set $\xi_1^{(i)} \gg \xi_0^{(i)} \forall j \in \mathcal{J}_{on}^+$ and $\xi_1^{(i)} \ll \xi_0^{(i)} \forall j \in \mathcal{J}_{off}^-$. This will be possible under two conditions: (1) there exists a unique i for each j, i.e. the controller sends a different signal to each device, (2) the equity constraint (8a) is not binding. If these conditions are met then the proposed method can reduce constraint violations within one time step as effectively as the benchmark. It should be noted that this robustness guarantee is with respect to the linear constraints; for voltage and branch bound constraints the accuracy of the robustness condition is limited by the accuracy of the power flow linearization.

V. Simulation Platform

We implemented the proposed method, alongside the benchmark with an agent-based formulation in Pyomo [38] – an optimization toolkit in Python that allows you to formulate optimization problems using various solvers. We have tested the platform using CPLEX [39], but the Pyomo framework allows for other mixed-integer linear programming (MILP) solvers to be used. MILP problems are a standard form and a variety of industry solvers are available which can be used to solve them. While the continuous relaxation (LP) are convex and will converge to a guaranteed global optimum, the mixed-integer variables introduce non-convexity. Therefore, solvers typically use a combination of branch-and-bound and heuristics to sequentially find a solution.

We are happy to release our simulation platform and all required data as open source, available at: https://github.com/constancecrozier/robustDR. The simulations in this paper were run on a MacBook Pro with a 2.0GHz quad-core Core i5 processor and 16GB of memory.

A. Computational Performance

Figure 4 shows the average time taken to compute price signals using a 5m resolution as the number of devices and time horizon grows. Deadline-driven devices are fast to converge; 100 devices with 288 time-steps solves within 1 second. Thermal devices scale worse with time horizon compared to deadline-driven devices. This is likely because for thermal devices the thermal leakage over multiple hours is more complex to calculate – whereas for deadline-driven devices there is no time-dependant loss term to calculate. This performance is consistent with the application because, although the problem is centralized, the number of devices is likely to remain small due to the nature of the problem (local distribution level constraints). Note that if a larger number of devices were present, a coarser time resolution (e.g. 30m instead of 5m) can counter-act the increased complexity. Furthermore, the shown results includes no parallel processing, which could accelerate the convergence.

VI. Case Study

The proposed method was testing using case studies of residential networks in the San Francisco area. Building electricity data at 15 min resolution was taken from [40]. Local temperature and global horizontal irradiance at 10 min resolution was taken from [41]. For real-time prices, we used the LMP prices from the California Independent System

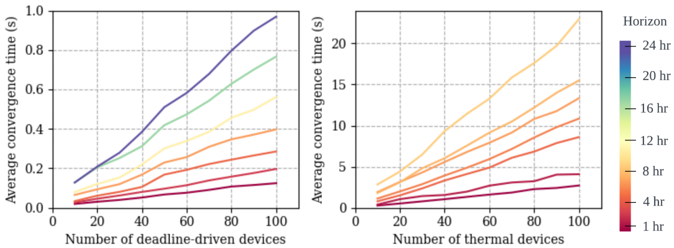

Fig. 4: Time taken to compute price signals for a varying number of devices and varying time horizons (at 5m resolution)

p	η	Δ_t	N_t	Θ	κ	M
7.0	0.9	$\frac{1}{12}$	288	30.0	10^6	10^6

TABLE I: The parameters used for Case Study 1

Operator (CAISO), for price forecasts we used the day-ahead system price. Both datasets can be accessed at [42].

A. Case Study 1: 30 Homes with EVs

We consider a network with 30 homes which all have EVs, and a transformer with a thermal rating of 30 kW. Charging data was obtained from the Electric Nation charging trial [4], which monitored the home charging over a multi-year period. Both vehicles and households were randomly selected from the available data. We assumed all chargers had a power rating of 7kW, and an efficiency of 90%. We used a planning horizon of 24 hours, and time intervals of five minutes. The parameters of the case study are summarized in Table I.

In order to evaluate the action of our algorithm, we consider three counterfactual scenarios to compare against. In the first case we consider where the EV charging is uncontrolled, meaning vehicles begin charging immediately when they are plugged in. The second case is top-down TSO controlled charging, where the EVs act to minimize their charging cost directly using the LMP prices. This case will give the lower bound on charging cost using the LMPs, but does not take into account local constraints. Finally, we consider direct control from an aggregator who aims to minimize LMP charging cost while respecting local constraints. Where all devices opt into demand response, this final case will provide a lower bound on the constraint violations which are possible with a rolling horizon approach. The aggregator will have direct control over all devices at once and perfect information about their end-use requirements. However, unlike the proposed approach, this method is not privacy-preserving.

First we consider the case where all EV owners select the ‘economy’ setting on their charger. We ran a simulation using demand data and prices starting from 2021/01/02 $\quad 0:00$ and continuing for a total of nine days. For controller estimates of the device deadlines and energy consumption, we applied Gaussian noise to the true values. The action of the proposed method compared to the three other cases is shown in Fig. 5. The top plot shows a subset of the total power demand, while the bottom plots show the charging costs and total constraint violations respectively. For simplicity, we visualize the summed violations rather than the individual constraint violations. Any non-zero value will imply accelerated degradation of network components which can be estimated using a device specific model (e.g. [43]).

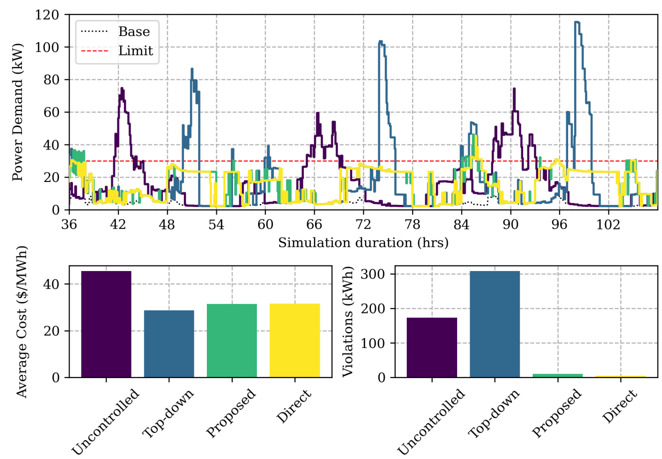

Fig. 5: A simulation with 30 EVs all opting into demand response under (a) uncontrolled charging, (b) TSO led charging, (c) the proposed method, and (d) direct control via an aggregator. The top plot shows the total power demand under each scheme for a subset of the simulation. The bottom left plot shows the average cost of charging (a proxy for TSO response) and the transformer violations under each scheme.

As expected, we see that the top-down method results in the lowest charging cost, while the direct method results in the lowest constraint violation. The top-down method results in by far the largest constraint violation – more than 50% higher than in the uncontrolled case. This demonstrates a classic case of the TSO-DSO conflict; by controlling charging we have sacrificed natural diversity in charging behavior, thus producing greater local constraint violation than if no action were taken. The proposed method of adjusting the local prices results in a 96% reduction in total constraint violations and a 8.4% increase in cost compared to the top-down case. This is compared to a 98% and 8.2% in the direct control case. Although this case satisfies the constraints derived in Section IV, constraint violations are slightly higher in the proposed case due to errors in device parameter estimates – as expected, additional violations are resolved within one time step.

Next, we consider the case where not all consumers opt in to DR. This is important because sometimes consumers may want to prioritize their consumption, for example if they are unsure when their EV will be next needed. Here we randomly assigned 50% of EVs to the ‘economy’ setting and 50% to the ‘priority’ setting. Figure 6 shows analysis of the same simulation with this single difference. In this case, we find that the proposed method actually results in lower constraint violations compared to the direct control case. This is because the aggregator has no visibility of the devices opting out of DR, so they coordinate the EVs as though there are no other smart devices on the system.

B. Case Study 2: 100 Homes with HVAC

Our second case study focused on homes with heating, ventilation and air-conditioning (HVAC) systems that can

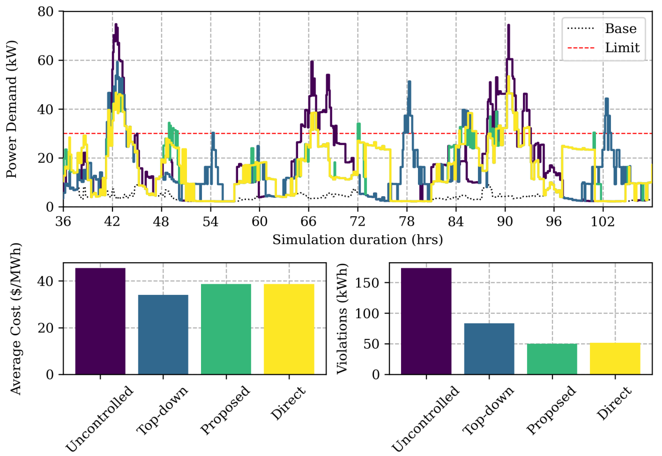

Fig. 6: A similar simulation to the one shown in Fig. 5, but with half of the devices opting out of demand response

p^C	η^C	p^H	η^H	r	C	k	Δ_t	N_t	Θ
0.35	-20	0.45	10	$\frac{1}{60}$	1.5×10^7	5	$\frac{1}{12}$	72	75

TABLE II: The parameters used for Case Study 2

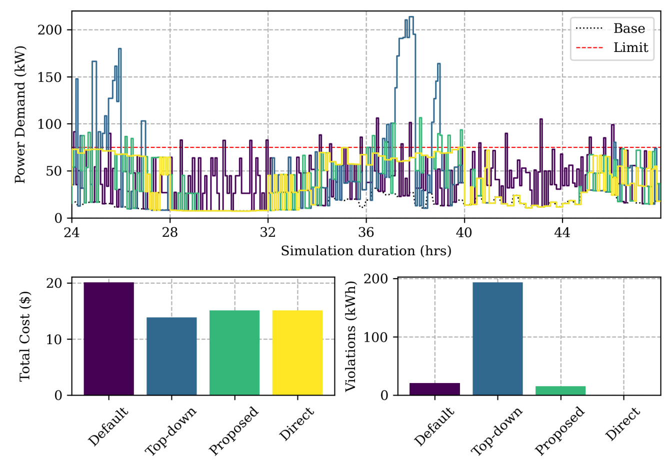

Fig. 7: The results for a case study from 100 homes with smart HVAC systems. Default: the devices act to minimize energy consumption, Top-down: the devices minimize energy cost using LMPs, Proposed: the proposed method where LMPs are altered to protect constraints, Direct: aggregator based control to minimize cost subject to network constraints.

participate in DR. We model the HVAC as two smart devices: one for heating and one for cooling. The building and device parameters are summarised in Table II, where the superscripts ^{H,C} are used to distinguish the heating and cooling devices respectively.

For our case study we assumed 10 nodes, each with 10 buildings with smart HVAC connected to them. This means that 10 devices receive the same signal from the controller, whereas in the previous case study there was only one device per grid node. We used temperature and demand data starting from $2020/01/02 \quad 0:00$ and ran a simulation of 48 hours length. Each building had the same thermal parameters, but different temperature bounds and initial temperatures – providing diversity between device requirements. The temperature bounds were randomly generated, with lower bounds ranging from 16°C to 21°C and upper bounds from 20°C to 24°C. The initial temperature for each building was randomly selected between the chosen bounds. The results are shown in Fig. 7.

As with the EV case study, we see that the top-down transmission pricing strategy results in worse violations than the default case. In fact, in this case the diversity between device requirements is sufficient to avoid almost all constraint violations in the default strategy. However, we can see from the direct control benchmark that a significant price reduction (and hence TSO response) is possible while respecting the distribution constraint. The proposed method achieves a near optimal TSO price response, with similar constraint violations as the default case. This is an impressive result, especially given that the controller is forced to send the same signal to ten devices at once. It is likely that the transformer constraint violations could be further reduced if different control signals were sent to each device.

C. Discussion

In order to observe the action of the controller, we consider the output and control signal sent to a single building with both an EV and HVAC system. Figure 8 shows the power demand, controller estimates for c, and control signals throughout a 16 hour simulation. The bottom plot shows the controller signal compared to the LMP prices. Note that both devices receive the same control signal as they are at the same grid location. It can be seen from the bottom plot that the controller chooses a significantly more volatile signal than the transmission prices. The controller signal generally exaggerates the LMP signal, but in some cases opposes it. This is likely done to coordinate the consumption of devices at this node with those in other grid locations. The middle plot shows the controllers estimate of c, the threshold price each device. Note that the controller does not have an estimate for the EV while it is disconnected. We can see that the EV is consistently more price sensitive, having a lower value at which it will turn on. This is reasonable, because the EV's energy requirement does not increase by shifting consumption, so the devices pay no penalty for shifting demand. The only parameters which must be chosen by the controller are the estimates of the device requirements. A sensitivity analysis revealed that for thermal devices the controller works relatively well with naive estimates of temperature bounds. However, for deadline-driven devices the performance drops significantly with poor estimates of the device's energy requirement. We noticed that all instances of significantly increased violations were in cases where $\hat{E}^{(j)}$ was an underestimate. This is likely because once the device has consumed the predicted energy requirement $\hat{E}^{(j)}$ the controller assumes that the device will no longer consumer power. This means the controller no longer attempts to control the device's output and does not account for its consumption in the distribution constraint. Conditional adjustments to estimates based on observed behavior, and purposefully conservative (over-)estimates could resolve this

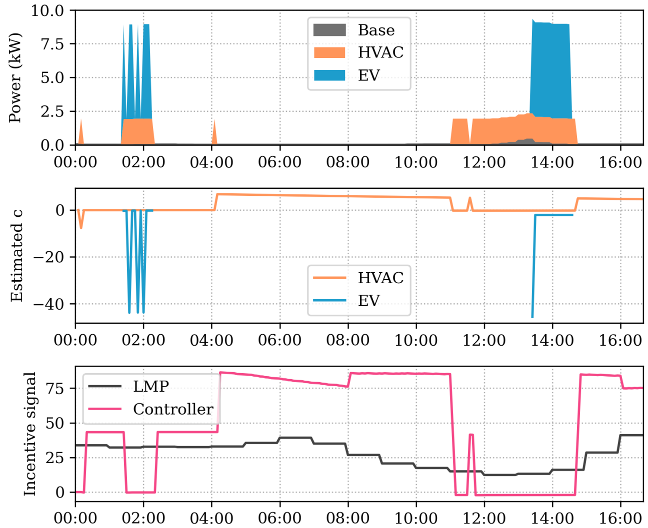

Fig. 8: Illustration of the action of the controller on a single building with both an EV and HVAC. Top: resulting power consumption, middle: controller estimated value of device point, bottom: original and adjusted price signal.

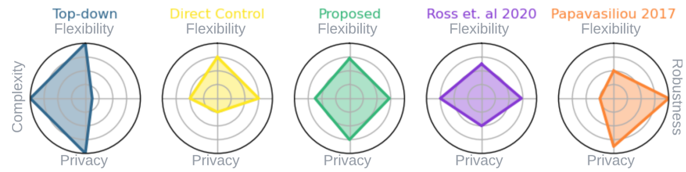

Fig. 9: A radar chart showing a comparison of several methods over: computational complexity, transmission level flexibility, distribution constraint robustness, and consumer privacy.

issue.

In Fig. 9 we compare the methods presented in this paper with two similar methods from the literature across four criteria: the amount of flexibility exploited, robustness to distribution constraints, consumer privacy, and computational cost. We consider the method proposed by Ross and Mathieu [12], which imposes constraints on aggregators in order to protect the distribution network. Additionally we consider distribution system marginal prices, first proposed by Papavasiliou [6], where the prices will reflect the constraints of the local distribution network. It can be seen that while the proposed method is not the highest scoring in any of these single criteria, it balances all objectives reasonably. The top-down method scores extremely well in flexibility exploitation, complexity, and privacy but offers no protection to distribution system constraints. The direct control benchmark achieves the theoretically optimal trade-off between flexibility and distribution constraint protection. However, it requires consumers to share sensitive information and the centralized formulation means the computational complexity grows poorly with number of devices. The method proposed in [12] protects distribution system constraints in a more computationally efficient manner, but still uses aggregators so consumers must still communicate device requirements. The locational pricing method protects both privacy and distribution constraints, but is extremely computationally intensive and does not directly maximize flexibility.

VII. Conclusion

In this paper we have developed a demand response (DR) framework for coordinating TSO and DSO objectives that preserves the privacy of residential consumers. Under this paradigm, the DSO controller generates a control signals to send to each distribution node which incentivize devices to align with TSO pricing, but cognizant of distribution constraints. We developed an open-source framework for testing this control scheme with a variety of flexible devices, including both electric vehicles and smart HVAC systems. We showed that if a unique signal is sent to each device and the equity constraint is not binding, constraint violations will be resolved within one time step. Using case studies, we demonstrate that the proposed method achieves close to the direct control benchmark in the case where all devices participate in demand response. Furthermore, we show that the proposed method can actually surpass the direct control benchmark when not all devices opt-in to DR.

Appendix

We start with the original equation for heat change (4). Then we assume the building is being held at it's temperature bound so $T_{t+1}^{(j)} = T_t^{(j)} = T_{min}^{(j)}$, giving:

$$0 = \Big(\frac{T_t^{out} - T_{min}^{(j)}}{R^{(j)}C^{(j)}} + \frac{k^{(j)}}{C^{(j)}}I_t - \frac{\eta^{(j)}}{C^{(j)}}x_t^{(j)}p^{(j)}10^3\Big)3600\Delta_t \tag{15}$$

Then we make the simplification that the irradiance term is zero, for heating this will result in a conservative estimate as the overall heat loss will be overestimated

$$0 = \Big(\frac{T_t^{out} - T_{min}^{(j)}}{R^{(j)}C^{(j)}} + \frac{\eta^{(j)}}{C^{(j)}}x_t^{(j)}p^{(j)}10^3\Big)3600\Delta_t \tag{16}$$

We can rearrange the equation to solve for power loss:

$$x_t^{(j)}p^{(j)} = \frac{T_{min}^{(j)} - T_t^{out}}{R^{(j)}\eta^{(j)}}10^{-3} \tag{17}$$

Finally, we inflate the cost by the energy loss so we multiply by the amount of time until the end of the horizon $\Delta_t(N_t - t)$.

References

[1] R. De Coninck and L. Helsen, "Quantification of flexibility in buildings by cost curves–methodology and application," *Applied Energy*, vol. 162, pp. 653–665, 2016.

[2] C. Crozier, T. Morstyn, and M. McCulloch, "The opportunity for smart charging to mitigate the impact of electric vehicles on transmission and distribution systems," *Applied Energy*, vol. 268, p. 114973, 2020.

[3] Octopus Energy, "Octopus agile: The 100% green electricity tariff with plunge pricing," 2018. [Online]. Available: https://octopus.energy/agile/

[4] Electric Nation, "Smart charging project: Customer trial final report," 2019. [Online]. Available: https://electricnation.org.uk/

[5] A. Vicente-Pastor, J. Nieto-Martin, D. W. Bunn, and A. Laur, "Evaluation of flexibility markets for retailer–dso–tso coordination," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 2003–2012, 2019.

[6] A. Papavasiliou, "Analysis of distribution locational marginal prices," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4872–4882, 2017.

[7] D. C. Matisoff, R. Beppler, G. Chan, and S. Carley, "A review of barriers in implementing dynamic electricity pricing to achieve cost-causality," *Environmental Research Letters*, vol. 15, no. 9, p. 093006, 2020.

[8] X. Zhou, C.-Y. Chang, A. Bernstein, C. Zhao, and L. Chen, "Economic dispatch with distributed energy resources: Co-optimization of transmission and distribution systems," *IEEE Control Systems Letters*, vol. 5, no. 6, pp. 1994–1999, 2021.

[9] Y. Li, M. Han, Z. Yang, and G. Li, "Coordinating flexible demand response and renewable uncertainties for scheduling of community integrated energy systems with an electric vehicle charging station: A bi-level approach," *IEEE Transactions on Sustainable Energy*, vol. 12, no. 4, pp. 2321–2331, 2021.

[10] S. M. Mohseni-Bonab, I. Kamwa, A. Moeini, and A. Rabiee, "Voltage security constrained stochastic programming model for day-ahead bess schedule in co-optimization of td systems," *IEEE Transactions on Sustainable Energy*, vol. 11, no. 1, pp. 391–404, 2020.

[11] N. I. Nimalsiri, C. P. Mediwaththe, E. L. Ratnam, M. Shaw, D. B. Smith, and S. K. Halgamuge, "A survey of algorithms for distributed charging control of electric vehicles in smart grid," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 11, 2020.

[12] S. C. Ross and J. L. Mathieu, "A method for ensuring a load aggregator's power deviations are safe for distribution networks," *Electric Power Systems Research*, vol. 189, p. 106781, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779620305848

[13] A. Karapetyan, S. K. Azman, and Z. Aung, "Assessing the privacy cost in centralized event-based demand response for microgrids," in *2017 IEEE Trustcom/BigDataSE/ICESS*, 2017, pp. 494–501.

[14] H. Yu, J. Zhang, J. Ma, C. Chen, G. Geng, and Q. Jiang, "Privacy-preserving demand response of aggregated residential load," *Applied Energy*, vol. 339, p. 121018, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261923003823

[15] E. J. Palacios-Garcia, X. Carpent, J. W. Bos, and G. Deconinck, "Efficient privacy-preserving aggregation for demand side management of residential loads," *Applied Energy*, vol. 328, p. 120112, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261922013691

[16] X. Kou, F. Li, J. Dong, M. Starke, J. Munk, Y. Xue, M. Olama, and H. Zandi, "A scalable and distributed algorithm for managing residential demand response programs using alternating direction method of multipliers (admm)," *IEEE Transactions on Smart Grid*, vol. 11, no. 6, pp. 4871–4882, 2020.

[17] J. Li, C. Li, Z. Wu, X. Wang, K. L. Teo, and C. Wu, "Sparsity-promoting distributed charging control for plug-in electric vehicles over distribution networks," *Applied Mathematical Modelling*, vol. 58, pp. 111–127, 2018.

[18] Y. Mou, H. Xing, Z. Lin, and M. Fu, "Decentralized optimal demand-side management for PHEV charging in a smart grid," *IEEE Transactions on Smart Grid*, vol. 6, no. 2, pp. 726–736, 2015.

[19] M. U. Hassan, M. H. Rehmani, and J. Chen, "Differentially private dynamic pricing for efficient demand response in smart grid," in *IEEE International Conference on Communications (ICC)*, 2020.

[20] F. Zobiri, M. Gama, S. Nikova, and G. Deconinck, "A privacy-preserving three-step demand response market using multi-party computation," in *2022 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, 2022, pp. 1–5.

[21] W. Chen, A. Zhou, P. Zhou, L. Gao, S. Ji, and D. Wu, "A privacy-preserving online learning approach for incentive-based demand response in smart grid," *IEEE Systems Journal*, vol. 13, no. 4, 2019.

[22] W. Fu, Y. Wan, J. Qin, Y. Kang, and L. Li, "Privacy-preserving optimal energy management for smart grid with cloud-edge computing," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, 2022.

[23] A. Ghasemkhani, L. Yang, and J. Zhang, "Learning-based demand response for privacy-preserving users," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 9, pp. 4988–4998, 2019.

[24] A. Pigott, C. Crozier, K. Baker, and Z. Nagy, "Gridlearn: Multiagent reinforcement learning for grid-aware building energy management," *Electric Power Systems Research*, vol. 213, p. 108521, 2022.

[25] A. Halder, X. Geng, P. Kumar, and L. Xie, "Architecture and algorithms for privacy preserving thermal inertial load management by a load serving entity," in *IEEE Power Energy Society General Meeting*, 2017.

[26] B. Moreno, A. J. Lopez, and M. T. García-Alvarez, "The electricity prices in the european union. the role of renewable energies and regulatory electric market reforms," *Energy*, vol. 48, no. 1, 2012.

[27] S. P. Burger, J. D. Jenkins, S. C. Huntington, and I. J. Perez-Arriaga, "Why distributed?: A critical review of the tradeoffs between centralized and decentralized resources," *IEEE Power and Energy Magazine*, vol. 17, no. 2, pp. 16–24, 2019.

[28] R. Li, Q. Wu, and S. S. Oren, "Distribution locational marginal pricing for optimal electric vehicle charging management," *IEEE Transactions on Power Systems*, vol. 29, no. 1, pp. 203–211, 2013.

[29] G. De Zotti, S. A. Pourmousavi, J. M. Morales, H. Madsen, and N. K. Poulsen, "Consumers' flexibility estimation at the tso level for balancing services," *IEEE Transactions on Power Systems*, vol. 34, no. 3, pp. 1918–1930, 2019.

[30] D. H. Nguyen, T. Narikiyo, and M. Kawanishi, "Optimal demand response and real-time pricing by a sequential distributed consensus-based admm approach," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4964–4974, 2018.

[31] S. Naderi, G. Pignatta, S. Heslop, I. MacGill, and D. Chen, "Demand response via pre-cooling and solar pre-cooling: A review," *Energy and Buildings*, vol. 272, p. 112340, 2022.

[32] A. Naeem, A. Shabbir, N. Ul Hassan, C. Yuen, A. Ahmad, and W. Tushar, "Understanding customer behavior in multi-tier demand response management program," *IEEE Access*, vol. 3, 2015.

[33] X. Kou, F. Li, J. Dong, M. Starke, J. Munk, T. Kuruganti, and H. Zandi, "A distributed energy management approach for residential demand response," in *ICSGSC*, 2019, pp. 170–175.

[34] S. Rouchier, M. Rabouille, and P. Oberlé, "Calibration of simplified building energy models for parameter estimation and forecasting: Stochastic versus deterministic modelling," *Building and Environment*, pp. 181–190, 2018.

[35] A. Vecchietti, S. Lee, and I. E. Grossmann, "Modeling of discrete/continuous optimization problems: characterization and formulation of disjunctions and their relaxations," *Computers & chemical engineering*, vol. 27, no. 3, pp. 433–448, 2003.

[36] C. Crozier, K. Baker, and B. Toomey, "Feasible region-based heuristics for optimal transmission switching," *Sustainable Energy, Grids and Networks*, vol. 30, p. 100628, 2022.

[37] A. Bernstein, C. Wang, E. Dall'Anese, J.-Y. Le Boudec, and C. Zhao, "Load flow in multiphase distribution networks: Existence, uniqueness, non-singularity and linear models," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 5832–5843, 2018.

[38] W. E. Hart, J.-P. Watson, and D. L. Woodruff, "Pyomo: modeling and solving mathematical programs in python," *Mathematical Programming Computation*, vol. 3, no. 3, pp. 219–260, 2011.

[39] Cplex, IBM ILOG, "V12. 1: User's manual for cplex," *International Business Machines Corporation*, vol. 46, no. 53, p. 157, 2009.

[40] E. Wilson et. al, "End-use load profiles for the u.s. building stock," 2021.

[41] National Renewable Energy Laboratory (NREL), "NSRDB: National solar radiation database," 2022. [Online]. Available: https://nsrdb.nrel.gov/

[42] California ISO , "Open access same-time information system (oasis)," 2022. [Online]. Available: http://oasis.caiso.com/mrioasis/logon.do

[43] B. Lesieutre, W. Hagman, and J. Kirtley, "An improved transformer top oil temperature model for use in an on-line monitoring and diagnostic system," *IEEE Trans. Power Deliv.*, vol. 12, no. 1, pp. 249–256, 1997.