

Optimal demand response: problem formulation and deterministic case

Lijun Chen, Na Li, Libin Jiang, and Steven H. Low

Abstract We consider a set of users served by a single load-serving entity (LSE). The LSE procures capacity a day ahead. When random renewable energy is realized at delivery time, it manages user load through real-time demand response and purchases balancing power on the spot market to meet the aggregate demand. Hence optimal supply procurement by the LSE and the consumption decisions by the users must be coordinated over two timescales, a day ahead and in real time, in the presence of supply uncertainty. Moreover, they must be computed jointly by the LSE and the users since the necessary information is distributed among them. In this paper we present a simple yet versatile user model and formulate the problem as a dynamic program that maximizes expected social welfare. When random renewable generation is absent, optimal demand response reduces to joint scheduling of the procurement and consumption decisions. In this case, we show that optimal prices exist that coordinate individual user decisions to maximize social welfare, and present a decentralized algorithm to optimally schedule a day in advance the LSE's procurement and the users' consumptions. The case with uncertain supply is reported in a companion paper.

1 Introduction

1.1 Motivation

There is a large literature on various forms of load side management from the classical direct load control to the more recent real-time pricing [1, 2]. Direct load control

Lijun Chen, Na Li, Libin Jiang and Steven H. Low
Engineering and Applied Science, California Institute of Technology, USA
e-mail: {chenlj, nali, libinj, slow}@caltech.edu

To appear: In Control and Optimization Theory for Electric Smart Grids. Aranya Chakraborty and Marija Ilic, Editors, Springer 2011

in particular has been practised for a long time and optimization methods have been proposed to minimize generation cost e.g. [3, 4, 5, 6], maximize utility's profit e.g. [7], or minimize deviation from users' desired consumptions e.g. [8, 9], sometimes integrated with unit commitment and economic dispatch e.g. [4, 10]. Almost all demand response programs today target large industrial or commercial users, or, in the case of residential users, a small number of them, for two, among other, important reasons. First, demand side management is invoked rarely to mostly cope with a large correlated demand spike due to weather or a supply shortfall due to faults, e.g., during a few hottest days in summer. Second, the lack of ubiquitous two-way communication in the current infrastructure prevents the participation of a large number of diverse users with heterogeneous and time-varying consumption requirements. Both reasons favor a simple and static mechanism involving a few large users that is sufficient to deal with the occasional need for load control, but both reasons are changing.

Renewable sources can fluctuate rapidly and by large amounts. As their penetration continues to grow, the need for regulation services and operating reserves will increase, e.g., [11, 12]. This can be provided by additional peaker units, at a higher cost, or supplemented by real-time demand response [13, 14, 15, 12, 16]. We believe that demand response will not only be invoked to shave peaks and shift load for economic benefits, but will increasingly be called upon to improve security and reduce reserves by adapting elastic loads to intermittent and random renewable generation [17]. Indeed, the authors of [12, 18, 19] advocate the creation of a distribution/retail market to encourage greater load side participation as an alternative source for fast reserves. Such application however will require a much faster and more dynamic demand response than practised today. This will be enabled in the coming decades by the large-scale deployment of a sensing, control, and two-way communication infrastructure, including the flexible AC transmission systems, the GPS-synchronized phasor measurement units, and the advanced metering infrastructure, that is currently underway around the world [20].

Demand response in such context must allow the participation of a large number of users, and be dynamic and distributed. Dynamic adaptation by hundreds of millions of end users on a sub-second control timescale, each contributing a tiny fraction of the overall traffic, is being practised everyday on the Internet in the form of congestion control. Even though both the grid and the Internet are massive distributed nonlinear feedback control systems, there are important differences in their engineering, economic, and regulatory structures. Nonetheless the precedence on the Internet lends hope to a much bigger scale and more dynamic and distributed demand response architecture and its benefit to grid operation. *Ultimately it will be cheaper to use photons than electrons to deal with a power shortage.* Our goal is to design algorithms for such a system.

1.2 Summary

Specifically we consider a set of users that are served by a single load-serving entity (LSE). The LSE may represent a regulated monopoly like most utility companies in the United States today, or a non-profit cooperative that serves a community of end users. Its purpose is (possibly regulated) to promote the overall system welfare. The LSE purchases electricity on the wholesale electricity markets (e.g., day-ahead, real-time balancing, and ancillary services) and sells it on the retail market to end users. It provides two important values: it aggregates loads so that the wholesale markets can operate efficiently, and it hides the complexity and uncertainty from the users, in terms of both power reliability and prices. Our model captures three important features:

- **Uncertainty.** Part of the electricity supply is from renewable sources such as wind and solar, and thus uncertain.
- **Supply and demand.** LSE's supply decisions and the users' consumption decisions must be jointly optimized.
- **Two timescale.** The LSE must procure capacity on the day-ahead wholesale market while user consumptions should be adapted in real time to mitigate supply uncertainty.

Hence the key is the coordination of day-ahead procurement and real-time demand response over two timescales in the presence of supply uncertainty. Moreover, the optimal decisions must be computed jointly by the LSE and the users as the necessary information is distributed among them. The goal of this paper is to formulate this problem precisely. Due to space limitation, we can only fully treat the case without supply uncertainty. Results for the case with supply uncertainty are summarized here, but fully developed in a companion paper [21].

Suppose each user has a set of appliances (electric vehicle, air conditioner, lighting, battery, etc.). She (or her energy management system) is to decide how much power she should consume in each period $t = 1, \dots, T$ of a day. The LSE needs to decide how much capacity it should procure a day ahead and, when the random renewable energy is realized at real time, how much balancing power to purchase on the spot market to meet the aggregate demand. In Section 2, we present our user and supply models, and formulate the overall problem as an $(1 + T)$ -period dynamic program to maximize expected social welfare. The key idea is to regard the LSE's day-ahead decision as the control in period 0 and the users' consumption decisions as controls in the subsequent periods $t = 1, \dots, T$. By unifying several models in the literature, our user model incorporates a large class of appliances. Yet, it is simple, thus analytically tractable, where each appliance is characterized by a utility function and a set of linear consumption constraints.

In Section 3, we consider the case without renewable generation. In the absence of uncertainty it becomes unnecessary to adapt user consumptions in real-time and hence supply and consumptions can be optimally scheduled at once instead of over two days. We show that optimal prices exist that coordinate individual users' decisions in a distributed manner, i.e., when users selfishly maximize their own surplus

under the optimal prices, their consumption decisions turn out to also maximize the social welfare. We develop an offline distributed algorithm that jointly schedules the LSE's procurement decisions and the users' consumption decisions for each period in the following day. The algorithm is decentralized where the LSE only knows the aggregate demand but not user utility functions or consumption constraints, and the users do not need to coordinate among themselves but only respond to the common prices from the LSE.

With renewable generation, the uncertainty precludes pure scheduling and calls for real-time consumption decisions that adapt to the realization of the random renewable generation. Moreover, this must be coordinated with procurement decisions over two timescales to maximize the expected welfare. Distributed algorithms for optimal demand response in this case and the impact of uncertainty on the optimal welfare are developed in the companion paper [21].

Finally we conclude in Section 4 with some limitations of this paper.

We make two remarks. First the effectiveness of real-time pricing for demand response is still in active research. On the one hand, empirical studies have shown consistently that price elasticity is low and heterogeneous; see [22, 23, 24] and references therein. On the other hand, there are strong economic arguments that real-time retail prices improve the efficiency of the overall system by allowing users to dynamically adapt their loads to shortages, with potential benefits far exceeding the cost of implementation [18]. Moreover, the long-run efficiency gain is likely to be significant even if demand elasticity is small, but unfortunately, the popular open-loop time-of-use pricing may capture a very small share of the efficiency gain of real-time pricing [25]. We neither argue for nor against real-time pricing. Indeed we do not consider in this paper the economic issues associated with such a system, such as locational marginal prices, revenue-adequacy, etc. What we refer to as 'prices' are simply control signals that provide the necessary information for users to adapt their consumption in a distributed, yet optimal, manner. Whether this control signal should be linked to monetary payments to provide the right incentive for demand response is beyond the scope of this paper, i.e., we do not address the important issue of how to incentivize users to respond to supply and demand fluctuations.¹

Second, unlike many current systems, the kind of large-scale distributed demand response system envisioned here must be fully automated. Human users set parameters that specify utility functions and consumption constraints and may change them on a slow timescale, but the algorithms proposed here will execute automatically and transparently to optimize social welfare. The traditional direct load control approach assumes that the controller (e.g. a utility company) knows the user consumption requirements, in the form of payback characteristics of the deferred load, and can optimally schedule deferred consumptions and their paybacks centrally. This is reasonable for the current system where the participating users are few and their requirements are relatively static. We take the view that the utilities and requirements of user consumptions are diverse and private. It is not practical, *nor necessary*, to have direct access to such information in order to optimally coordinate their con-

¹ See however [19] for a discussion on some implementation issues of real-time pricing for retail markets and a proposal for the Italian market.

somptions in a large, distributed, and dynamic system of the future. The algorithm presented here is an example that can achieve optimality without requiring users to disclose their private information.

1.3 Other related work

A large literature exists on demand response. Besides those cited above, more recent works include, e.g., [26, 27] on load control of thermal mass in buildings, [28, 29, 30] on residential load control through coordinated scheduling of different appliances, [31, 32, 33] on the scheduling of plug-in electric vehicle charging, and [34] on the optimal allocation of a supply deficit (rationing) among users using their supply functions. Load side management in the presence of uncertain supply has also been considered in [16, 10, 35, 36, 12, 37]. Unlike the conventional approach that compensates for the uncertainty to create reliable power, the authors of [16] advocate selling interruptible power and designs service contracts, based on [38], that can achieve greater efficiency than the conventional approach. In [10] various optimization problems are formulated that integrate demand response with economic dispatch with ramping constraints and forecasts of renewable power and load. Both centralized dispatch using model predictive control and decentralized dispatch using prices, or supply and demand functions, are considered. A two-period stochastic dispatch model is studied in [35] and a settlement scheme is proposed that is revenue-adequate even in the presence of uncertain supply and demand. A queueing model is analyzed in [36] where the queue holds deferrable loads that arise from random supply and demand processes. Conventional generation can be purchased to keep the queue small and strategies are studied to minimize the time-average cost. The models that are closest to ours, developed independently, are [12, 37]. All our models include random renewable generation, consider both day-ahead and real-time markets, and allow demand response, but our objectives and system operations are quite different. The authors of [12] advocate the establishment of a retail market where users (e.g., PHEVs) can buy power from or sell reserves, in the form of demand response capability, to their LSE. The paper formulates the LSE's and users' problems as dynamic programs that minimize their expected costs over their bids, which can be either simple, uncorrelated (price, quantity) pairs for each period, or complex, (price, quantity) pairs with temporal correlations. The model in [37] includes non-elastic users that are price non-responsive, and elastic users that can either leave the system or defer their consumptions when the electricity price is high. The goal is to maximize LSE's profit over day-ahead procurement, day-ahead prices for non-elastic users, and real-time prices for elastic users.

1.4 Notations

Given quantities such as the demands $q_{ia}(t)$ from appliance a of user i in period t , $q_{ia} := (q_{ia}(t), t \in \mathcal{T})$ denotes the vector of demands at different times, $q_i(t) := (q_{ia}(t), a \in \mathcal{A}_i)$ the vector of demands of different appliances, $q_i := (q_i(t), t \in \mathcal{T})$ the vector of demands of i 's appliances at different times, and $q := (q_i, \forall i)$ the vector of all demands. Similarly for the aggregate demands $Q_i(t) = \sum_{a \in \mathcal{A}_i} q_{ia}(t)$, $Q_{ia} := \sum_i q_{ia}(t)$, Q_i , Q , etc. Script letters denote sets, e.g., \mathcal{N} , \mathcal{A}_i , \mathcal{T} . Small letters denote individual quantities, e.g., $q_{ia}(t)$, q_{ia} , $q_i(t)$, q_i , q , etc. Capital letters denote aggregate quantities, e.g., $Q_i(t)$, Q_{ia} , $P_d(t)$, $P_r(t)$, $P_o(t)$, $P_b(t)$, etc. We use $q_{ia}(t)$, q_{ia} , $Q_i(t)$, etc. for loads and $P_d(t)$, $P_r(t)$, etc. for supplies. We sometimes write $\sum_i \sum_{a \in \mathcal{A}_i} q_{ia}(t)$ as $\sum_{i,a} q_{ia}(t)$. For any real a, b, c , $[a]_+ := \max\{a, 0\}$ and $[a]_b^c := \max\{b, \min\{a, c\}\}$. Finally, we write a vector as $x = (x_i, \forall i)$ without specifying whether it is a column or row vector so we can ignore the transpose sign to simplify the notation; the meaning should be clear from the context.

2 Model and problem formulation

Consider a set \mathcal{N} of N users that are served by a single load-serving entity (LSE). We use a discrete-time model with a finite horizon that models a day. Each day is divided into T periods of equal duration, indexed by $t \in \mathcal{T} := \{1, 2, \dots, T\}$. The duration of a period can be 5, 15, or 60 mins, corresponding to the time resolution at which energy dispatch or demand response decisions are made.

2.1 User model

Each user $i \in \mathcal{N}$ operates a set \mathcal{A}_i of appliances such as HVAC (heat, ventilation, air conditioner), refrigerator, and plug-in hybrid electric vehicle. User i may also possess a battery which provides further flexibility for optimizing its electricity consumption across time.

Appliance model. For each appliance $a \in \mathcal{A}_i$ of user i , $q_{ia}(t)$ denotes its energy consumption in period $t \in \mathcal{T}$, and q_{ia} the vector $(q_{ia}(t), \forall t)$ over the whole day. An appliance a is characterized by:

- a utility function $U_{ia}(q_{ia})$ that quantifies the utility user i obtains from using appliance a ;
- a $K_{ia} \times T$ matrix A_{ia} and a K_{ia} -vector η_{ia} such that the vector of power q_{ia} satisfies the linear inequality

$$A_{ia}q_{ia} \leq \eta_{ia}. \quad (1)$$

In general U_{ia} depends on the vector q_{ia} . In this paper, however, we consider four types of appliances whose utility functions take one of three simple forms. These

models are summarized in Table 1 and justified in detail in Appendix A. The utility of a type 1 or type 2 appliance is additive in t :²

$$U_{ia}(q_{ia}) := \sum_t U_{ia}(q_{ia}(t), t). \quad (2)$$

The utility of a type 3 appliance depends only on the aggregate consumption:

$$U_{ia}(q_{ia}) := U_{ia}\left(\sum_t q_{ia}(t)\right). \quad (3)$$

The utility of a type 4 appliance depends on the internal temperature and power consumptions in the past. It is of the form:

$$U_{iq}(q_{ia}) := \sum_t U_{ia}\left(T_{ia}(t) + \beta \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} q_{ia}(\tau)\right) \quad (4)$$

where $T_{ia}(t)$ is a given sequence of temperatures defined in equation (30) in Appendix A and α, β are given thermal constants. All utility functions are assumed to be continuously differentiable and concave functions for each t .

For example, some of our simulations in [39, 21] use the following time independent and additive utility function of form (2): let $y_{ia}(t)$ be a desired energy consumption by appliance a in period t ; then the function

$$U_{ia}(q_{ia}(t), t) := U_{ia}(q_{ia}(t)) := -(q_{ia}(t) - y_{ia}(t))^2 \quad (5)$$

measures the utility of following the desired consumption profile $y_{ia}(t)$. Such utility functions minimize user discomfort as advocated in [8, 9].

Table 1: Structure of utility functions and consumption constraints for appliances.

Appliances	Utility function	Consumption constraints	Examples
Type 1	(2)	(6)	Lightings
Type 2	(2)	(6), (7)	TV, video game, computer
Type 3	(3)	(6), (7)	PHEV, washers
Type 4	(4)	(6), (8)	HVAC, refrigerator
Battery	$-D_i(r_i)$	(6), (7)	$r_i = q_{ia}$ for battery a

The consumption constraints (1) for these appliances take three particular forms. First, for all appliances, the (real) power consumption must lie between a lower and an upper bound, possibly time-dependent:

$$\underline{q}_{ia}(t) \leq q_{ia}(t) \leq \bar{q}_{ia}(t). \quad (6)$$

² We abuse notation to use U_{ia} to denote both a function of vector q_{ia} and that of a scalar $q_{ia}(t)$; the meaning should be clear from the context.

An important character of an appliance is its allowable time of operation; e.g., an EV can be charged only between 9pm and 6am, TV may be on only between 7–9am and 6–12pm. If an appliance operates only in a subset $\mathcal{T}_{ia} \subseteq \mathcal{T}$ of periods, we require that $q_{ia}(t) = \bar{q}_{ia}(t) = 0$ for $t \notin \mathcal{T}_{ia}$ and $U_{ia}(0) = 0$. We therefore do not specify \mathcal{T}_{ia} explicitly in the description of utility functions and always sum over all $t \in \mathcal{T}$. The second kind of constraint specifies the range in which the aggregate consumption must lie:

$$\underline{Q}_{ia} \leq \sum_t q_{ia}(t) \leq \bar{Q}_{ia}. \quad (7)$$

The last kind of constraint is slightly more general (see derivation in Appendix A):

$$\underline{\eta}_{ia} \leq A_{ia} q_{ia} \leq \bar{\eta}_{ia}. \quad (8)$$

Battery model. We denote by B_i the battery capacity, by $b_i(t)$ the state of charge in period t , and by $r_i(t)$ the power (energy per period) charged to (when $r_i(t) \geq 0$) or discharged from (when $r_i(t) < 0$) the battery in period t . We use a simplified model of battery that ignores power leakage and other inefficiencies, where the state of charge is given by

$$b_i(t) = \sum_{\tau=1}^t r_i(\tau) + b_i(0). \quad (9)$$

The battery has an upper bound on charge rate, denoted by \bar{r}_i , and an upper bound on discharge rate, denoted by $-r_i$. We thus have the following constraints on $b_i(t)$ and $r_i(t)$:

$$0 \leq b_i(t) \leq B_i, \quad -r_i \leq r_i(t) \leq \bar{r}_i. \quad (10)$$

We assume any battery discharge is consumed by other appliances (zero leakage), and hence it cannot be more than what the appliances need:

$$-r_i(t) \leq \sum_{a \in \mathcal{A}_i} q_{ia}(t). \quad (11)$$

Finally, we impose a minimum on the energy level at the end of the control horizon: $b(T) \geq \gamma_i B_i$ where $\gamma_i \in [0, 1]$.

The cost of operating the battery is modeled by a function $D_i(r_i)$ that depends on the vector of charged/discharged power $r_i := (r_i(t), \forall t)$. This cost may correspond to the amortized purchase and maintenance cost of the battery over its lifetime, and depends on how fast/much/often it is charged and discharged; see an example $D_i(r_i)$ in [39]. The cost function D_i is assumed to be a convex function of the vector r_i .

Note that in this model, a battery is equivalent to an appliance: its utility function is $-D_i(r_i)$ and its consumption constraints (9), (10), and $b(T) \geq \gamma_i B_i$ are of the same form as (6)–(7) with $q_{ia} = r_i$. Therefore a battery can be specified simply as another appliance, in which case the constraint (11) requires that i 's aggregate demand be nonnegative, $\sum_{a \in \mathcal{A}_i} q_{ia}(t) + r_i(t) \geq 0$. This is summarized in Table 1. Henceforth,

we will often use appliances to also include battery and may not refer to battery explicitly when this does not cause confusion.

2.2 Supply model

We now describe a simple model of the electricity markets. The LSE procures power for delivery in each period t , in two steps. First it procures day-ahead capacities $P_d(t)$ for each period t a day in advance and pays for the capacity costs $c_d(P_d(t);t)$. The renewable power in each period t is a nonnegative random variable $P_r(t)$ and it costs $c_r(P_r(t);t)$. It is desirable to use as much renewable power as possible, for instance, if the renewable generation is owned by the LSE. For notational simplicity only, we assume $c_r(P;t) \equiv 0$ for all $P \geq 0$ and all t . Then at time t^- (real time), the random variable $P_r(t)$ is realized and used to satisfy demand. The LSE satisfies any excess demand by some or all of the day-ahead capacity $P_d(t)$ procured in advance and/or by purchasing balancing power on the real-time market. Let $P_o(t)$ denote the amount of the day-ahead power that the LSE actually uses and $c_o(P_o(t);t)$ its cost. Let $P_b(t)$ be the real-time balancing power and $c_b(P_b(t);t)$ its cost.

These real-time decisions $(P_o(t), P_b(t))$ are made by the LSE so as to minimize its total cost, as follows. Given the demand vector $q(t) := (q_{ia}(t), a \in \mathcal{A}_i, \forall i)$, let $Q(t) := \sum_{i,a} q_{ia}(t)$ be the total demand and $\Delta(Q(t)) := Q(t) - P_r(t)$ the excess demand, in excess of the renewable generation $P_r(t)$. Note that $\Delta(Q(t))$ is a random variable in and before period $t-1$, but its realization is known to the LSE at time t^- . Given excess demand $\Delta(Q(t))$ and day-ahead capacity $P_d(t)$, the LSE chooses $(P_o(t), P_b(t))$ that minimizes its total real-time cost, i.e., it chooses $(P_o^*(t), P_b^*(t))$ that solves the problem:

$$c_s(\Delta(Q(t)), P_d(t); t) := \min_{P_o(t), P_b(t)} \{ c_o(P_o(t); t) + c_b(P_b(t); t) \mid P_b(t) \geq 0, \\ P_o(t) + P_b(t) \geq \Delta(Q(t)), P_d(t) \geq P_o(t) \geq 0 \}. \quad (12)$$

Clearly $P_o^*(t) + P_b^*(t) = \Delta(Q(t))$ unless $\Delta(Q(t)) < 0$. The total cost is

$$c(Q(t), P_d(t); P_r(t), t) := c_d(P_d(t); t) + c_s(\Delta(Q(t)), P_d(t); t). \quad (13)$$

with $\Delta(Q(t)) := Q(t) - P_r(t)$. We assume that, for each t , $c_d(\cdot; t)$, $c_o(\cdot; t)$ and $c_b(\cdot; t)$ are increasing, convex, and continuously differentiable with $c_d(0; t) = c_o(0; t) = c_b(0; t) = 0$.

Example: supply cost

Suppose $c_b'(0) > c_o'(P), \forall P \geq 0$, i.e., the marginal cost of balancing power is strictly higher than the marginal cost of day-ahead power, the LSE will use the balancing power only after the day-ahead power is exhausted, i.e., $P_b(t) > 0$ only if $\Delta(Q(t)) > P_d(t)$. The solution $c_s(\Delta(Q(t)), P_d(t); t)$ of (12) in this case is particularly simple and (13) can be written explicitly in terms of c_b, c_o, c_b :

$$c(Q(t), P_d(t); P_r(t), t) = c_d(P_d(t); t) + c_o \left([\Delta(Q(t))]_0^{P_d(t)}; t \right) + c_b ([\Delta(Q(t)) - P_d(t)]_+; t). \quad (14)$$

i.e., the total cost consists of the capacity cost c_d and the energy cost c_o of day-ahead power, and the cost c_b of the real-time balancing power. \square

2.3 Problem formulation: welfare maximization

Recall that $q := (q(t), t \in \mathcal{T})$ and $Q(t) := \sum_{i,a} q_{ia}(t)$. The social welfare is the standard user utility minus supply cost:

$$W(q, P_d; P_r) := \sum_{i,a} U_{ia}(q_{ia}) - \sum_{t=1}^T c(Q(t), P_d(t); P_r(t), t). \quad (15)$$

As mentioned above the LSE's objective is not to maximize its profit through selling electricity, but rather to maximize the expected social welfare. Given the day-ahead decision P_d , the real-time procurement $(P_o(t), P_b(t))$ is determined by the simple optimization (13). This is most transparent in (14) for the special case: the optimal decision is to use day-ahead power $P_o^*(t)$ to satisfy any excess demand $\Delta(Q(t))$ up to $P_d(t)$, and then purchase real-time balancing power $P_b^*(t) = [\Delta(Q(t)) - P_d(t)]_+$ if necessary. Hence the maximization of (15) reduces to optimizing over day-ahead procurement P_d and real-time consumption q in the presence of random renewable generation $P_r(t)$. It is therefore critical that, in the presence of uncertainty, $q(t)$ should be decided after $P_r(t)$ have been realized at times t^- . P_d however must be decided a day ahead before $P_r(t)$ are realized.

The traditional dynamic programming model requires that the objective function be separable in time t . The welfare function in (15) is not as the first term $U_{ia}(q_{ia})$ depends on the entire control sequence $q_{ia} = (q_{ia}(t), \forall t)$. So does the consumption constraint (1). We now introduce an equivalent state space formulation of that will allow us to state precisely the overall optimization problem as an $(1+T)$ -period dynamic program.

Consider a dynamical system over an extended time horizon $t = 0, 1, \dots, T$. The control inputs are the LSE's day-ahead decision $P_d := (P_d(t), \forall t)$ in period 0 and the user's decisions $q(t)$ in each subsequent period. Let $v(t)$ denote the inputs, i.e., $v(0) = P_d$ and $v(t) = q(t)$, $t = 1, \dots, T$. Note that $v(0) \in \mathfrak{R}_+^T$ whereas $q(t) \in \mathfrak{R}^M$ where $M := \sum_{i=1}^N |\mathcal{A}_i|$. The system state $x(t) := (x^1(t), x_{ia}^2(t), x_{ia}^3(t), x_{ia}^4(t))$, $a \in \mathcal{A}_i, \forall i$ has four components, defined as follows:

- Without loss of generality, $x(0)$ starts from the origin.
- $x^1(t) \in \mathfrak{R}^T$ keeps track of the day-ahead decisions P_d : for each $t = 1, \dots, T$, $x^1(t) = P_d = (P_d(\tau), \tau = 1, \dots, T)$.

- $x_{ia}^2(t) \in \mathfrak{R}^{k_{ia}}$ of appropriate dimension k_{ia} for each (i, a) pair keeps track of the consumption constraint (1). The state definition and its transition are problem specific; see a concrete example in Section 2.4.
- $x^3(t) \in \mathfrak{R}_+$ keeps track of the random renewable power $x^3(0) = 0$, $x^3(t) = P_r(t)$, $t = 1, \dots, T$. The purpose of this state definition is merely notational, so that the control policy can depend on the *realization* of the random renewable power $P_r(t)$ through its dependence on state $x^3(t)$.
- $x_{ia}^4(t) \in \mathfrak{R}^{T-1}$ for each (i, a) pair tracks the user decisions $v_{ia}(t-1) = q_{ia}(t-1)$ in the previous period: $x_{ia}^4(1) = 0_{T-1}$, the $T-1$ dimensional zero vector; for each $t = 2, \dots, T$, the $(t-1)$ th component $[x_{ia}^4(t)]_{t-1}$ of $x_{ia}^4(t)$ is set to be the input $v_{ia}(t-1)$ and all the other components $[x_{ia}^4(t)]_\tau$ of $x_{ia}^4(t)$ remain the same as those of $x_{ia}^4(t-1)$, so that the final state $x_{ia}^4(T)$ is the vector $(q_{ia}(t), t = 1, \dots, T-1)$ of inputs up to period $T-1$. The first term in (15) is then a function of the state and input in period T , $U_{ia}(q_{ia}) = U_{ia}(x_{ia}^4(T), v_{ia}(T))$. This allows us to rewrite the welfare function in (15) in a form that is separable in t ; see below.

The above discussion is summarized by a time-varying state transition function f_t :

$$x(t+1) = f_t(x(t), v(t), P_r(t+1)), \quad t = 0, \dots, T$$

i.e., the new state $x(t+1)$ depends on the current state $x(t)$, the input $v(t)$, and the new random variable $P_r(t)$, and is therefore random. The consumption constraints (1), which may include the battery constraints, generally translate into constraints on the state $x^2(t)$ and input $v(t)$ and we represent this by $x(t) \in \mathcal{X}(t)$ and $v(t) \in \mathcal{V}(t) \subseteq \mathfrak{R}^M$, $M := \sum_{i=1}^N |\mathcal{A}_i|$. Sometimes these constraints also give rise to a terminal reward that we denote by $W_{T+1}(x(T+1))$.

Consider the class of feedback control laws $v(t) = \phi_t(x(t))$, where $\phi_0 : \mathcal{X}(0) \rightarrow \mathfrak{R}_+^T$ specifies the day-ahead decision P_d and $\phi_t : \mathcal{X}(t) \rightarrow \mathcal{V}(t)$ specifies the user decisions $q(t)$ for each period $t = 1, \dots, T$. Hence the control $v(t)$ depends only on the current state $x(t)$. Under the control law $\phi := (\phi_t, t = 0, \dots, T)$, the state evolves (stochastically) according to

$$x(t+1) = f_t(x(t), \phi_t(x(t)), P_r(t+1)). \quad (16)$$

We emphasize that $x(t)$ is obtained under policy ϕ even though this may not be explicit in the notation.

To make the welfare function in (15) separable in t , use (13) to define the welfare in each period t , under the control law ϕ , as a function of the current state $x(t)$ and the current input $v(t) = \phi_t(x(t))$:

$$W_t^\phi := W_t^\phi(x(t), v(t)) := \begin{cases} -\sum_{\tau=1}^T c_d([v(0)]_\tau; \tau), & t = 0 \\ -c_s(\Delta(Q^\phi(t)), [x^1(t)]_t; t), & 1 \leq t < T \\ \sum_{i,a} U_{ia}((x_{ia}^4(T), v_{ia}(T))) - c_s(\Delta(Q^\phi(T)), [x^1(T)]_T; T), & t = T \end{cases} \quad (17)$$

where $Q^\phi(t) = \sum_{i,a} [v(t)]_{ia}$ is the aggregate demand in period t under ϕ , and $v_{ia}(T) = q_{ia}(T)$ are the real-time consumption decisions in the last control period T . Then the welfare function in (15) is equivalent to

$$J^\phi := \sum_{t=0}^T W_t^\phi(x(t), v(t)) + W_{T+1}^\phi(x(T+1))$$

where the definition of the terminal reward $W_{T+1}^\phi(x(T+1))$ is problem specific. We can now state precisely our objective as the constrained maximization of the expected welfare over the control law ϕ :

$$\max_{\phi} \quad EJ^\phi = E \left(\sum_{t=0}^T W_t^\phi + W_{T+1}^\phi \right) \quad \text{s. t. } x^\phi(t) \in \mathcal{X}(t). \quad (18)$$

where the expectation is taken over $P_r(t), t = 1, \dots, T$.

Remark. An important assumption in this formulation is that the consumption constraints (1) can be modeled by an appropriate definition of states $x_{ia}^2(t)$, their transitions f_t , the constraint sets $\mathcal{X}(t), \mathcal{Y}(t)$, and possibly a terminal reward $W_{T+1}^\phi(x(T+1))$.

We now illustrate the problem formulation using a concrete example.

2.4 Example

To simplify the notation we make two assumptions that do not cause any loss of generality. First we use the total cost function c in (14) in the definition of the welfare function (15). Second we assume each user i has a single type-2 appliance and no battery (so we drop the subscript a). From Table 1, user utility functions are additive in time, $U_i(q_i) = \sum_t U_i(q_i(t); t)$ and the consumption constraints are

$$\underline{q}_i(t) \leq q_i(t) \leq \bar{q}_i(t), \quad \forall i \quad (19)$$

$$\bar{Q}_i \leq \sum_{t=1}^T q_i(t). \quad (20)$$

Since the utility functions are separable in t , we do not need to define $x^4(t)$. We now describe the $(1+T)$ -period dynamic program by specifying the definition of $x^2(t)$, the state transition function f_t , and the constraint sets $\mathcal{X}(t), \mathcal{Y}(t)$.

The system state $x(t) := (x^1(t), x^2(t), x^3(t))$ consists of three components of appropriate dimensions with

$$x(t) = (P_d, x^2(t), P_r(t)), \quad t = 1, \dots, T$$

where $x^2(t)$ is determined by the constraint (20). To simplify exposition, we make the important assumption that $P_r(t)$ are independent for different t ; see [21] for

a model without this independence assumption. Define $x_i^2(t)$ to be the remaining demand of user i at the beginning of each period t : $x_i^2(1) = \bar{Q}_i$, and for each $t = 1, \dots, T$, $x_i^2(t+1) = x_i^2(t) - v_i(t)$ where $v_i(t) = q_i(t)$. To enforce that $x^2(T+1) \leq 0$, we define the terminal cost $c_{T+1}(x(T+1)) = 0$ if $x^2(T+1) \leq 0_N$ and $c_{T+1}(x(T+1)) = \infty$ otherwise, where 0_n is the n -dimensional zero vector. Let the initial state be $x(0) = 0_{T+N+1}$. Denote $\bar{Q} := (\bar{Q}_i, \forall i)$. The system dynamics is then linear time-varying:

$$x(1) = x(0) + \begin{pmatrix} I_T \\ 0_{(N+1) \times T} \end{pmatrix} v(0) + \begin{pmatrix} 0_T \\ \bar{Q} \\ P_r(1) \end{pmatrix}$$

$$x(t+1) = \begin{pmatrix} I_{T+N} & 0_{T+N} \\ 0_{T+N} & 0 \end{pmatrix} x(t) - \begin{pmatrix} 0_{T \times N} \\ I_N \\ 0 \end{pmatrix} v(t) + \begin{pmatrix} 0_{T+N} \\ \mathbf{1} \end{pmatrix} P_r(t+1), \quad \forall 1 \leq t \leq T$$

where I_n is the $n \times n$ identify matrix, $0_{m \times n}$ the $m \times n$ zero matrix, and $P_r(T+1) := 0$.

The welfare in each period, under input sequence v , is (using (14))

$$W_0^v(x(0), v(0)) := - \sum_{\tau=1}^T c_d(P_d(\tau); \tau) = - \sum_{\tau=1}^T c_d([v(0)]_\tau; \tau)$$

and for $t = 1, \dots, T$,

$$W_t^v(x(t), v(t))$$

$$:= \sum_i U_i(q_i(t); t) - c_o \left([Q(t) - P_r(t)]_0^{P_d(t)}; t \right) - c_b \left([Q(t) - P_r(t) - P_d(t)]_+; t \right)$$

$$= \sum_i U_i(v_i(t); t) - c_o \left([\mathbf{1}v(t) - x^3(t)]_0^{[x^1(t)]_t}; t \right) - c_b \left([\mathbf{1}v(t) - x^3(t) - [x^1(t)]_t]_+; t \right)$$

where $\mathbf{1}$ is the (row) vector of 1's.

The constraint (19) yields the input constraint sets $\mathcal{V}(0) := \mathfrak{R}_+^T$ and, for $t = 1, \dots, T$, $\mathcal{V}(t) := \{q(t) \in \mathfrak{R}^N | \underline{q}(t) \leq q(t) \leq \bar{q}(t)\}$. There is no constraint on the state, i.e., $\mathcal{X}(t) = \mathfrak{R}^{T+N+1}$. Let $\phi := \{\phi_0 : \mathfrak{R}^{T+N+1} \rightarrow \mathfrak{R}_+^T, \phi_t : \mathfrak{R}^{T+N+1} \rightarrow \mathcal{V}(t), t = 1, \dots, T\}$ be the control policy so that $v(t) = \phi_t(x(t))$, $0 \leq t \leq T$. Then the welfare maximization problem (18) is

$$\max_{\phi} E \left(W_0^\phi(x(0), v(0)) + \sum_{t=1}^T W_t^\phi(x(t), v(t)) - c_{T+1}(x(T+1)) \right) \quad (21)$$

where the state $x(t)$ and the input $v(t)$ are obtained under policy ϕ .

In [21] we study the case with supply uncertainty in detail. We propose a distributed heuristic algorithm to solve the $(1+T)$ -period dynamic program. We prove that the algorithm is optimal when the welfare is quadratic and the LSEs procurement decisions are strictly positive. Otherwise, we bound the gap between the welfare achieved by the heuristic algorithm and the maximum. Simulation results sug-

gest that the performance of the heuristic algorithm is very close to optimal. As we scale up the size of a renewable generation plant, both its mean production and its variance will likely increase. As expected, the maximum welfare increases with the mean production, when the variance is fixed, and decreases with the variance, when the mean is fixed. More interesting, we prove that as we scale the size of the plant up, the maximum welfare increases.

3 Optimal scheduling without supply uncertainty

In this paper we only fully treat the case where there is no supply uncertainty, i.e., $P_r(t) \equiv 0$. Our goal is to optimally coordinate supply and demand to maximize social welfare. In the absence of uncertainty (our model also ignores demand uncertainty), it becomes unnecessary to adapt user consumptions in real-time and hence supply and consumptions can be optimally scheduled at once instead of over two days. Welfare maximization (18) then takes a simpler form and we develop an offline distributed algorithm that jointly optimizes the LSE's procurements and the users' consumptions for each period in the following day.

3.1 Optimal procurements and consumptions

We first consider LSE's procurement decisions. Recall that $Q_i(t) := \sum_{a \in \mathcal{A}_i} q_{ia}(t)$ and $\sum_i Q_i(t)$ is the aggregate demand in period t . With supply uncertainty, while P_d is decided a day ahead, the optimization (12) must be carried out in real time after $P_r(t)$ has been realized to obtain optimal $P_o(t), P_b(t)$. Here, on the other hand, all three decisions $(P_d(t), P_o(t), P_b(t))$ can be computed in advance in the absence of uncertainty. Hence, given an aggregate demand $\sum_i Q_i(t)$, the LSE solves (instead of (12)–(13)):

$$c \left(\sum_i Q_i(t); t \right) := \min_{P_d(t), P_o(t), P_b(t)} c_d(P_d(t); t) + c_o(P_o(t); t) + c_b(P_b(t); t) \quad (22)$$

$$\text{s. t. } P_o(t) + P_b(t) \geq \sum_i Q_i(t), \quad P_d(t) \geq P_o(t) \geq 0, \quad P_b(t) \geq 0$$

to obtain the total cost. The solution of (22) specifies the optimal decisions $(P_d^*(t), P_o^*(t), P_b^*(t))$ to satisfy the aggregate demand $\sum_i Q_i(t)$ for each period t in the following day.

It is not difficult to show that $c(\cdot, t)$ is a non-decreasing, convex, and continuously differentiable function for each t , so the problem (22) is convex. Since $c'_d(P; t) > 0$, the KKT condition implies that $P_d^*(t) = P_o^*(t)$ at optimality, i.e., it is optimal to exhaust all the day-ahead capacity. This is always possible because all procurement decisions are computed jointly without uncertainty. If we further assume that the marginal cost of the balancing power is higher than that of the day-

ahead power, i.e., $c'_b(0;t) > c'_d(P;t) + c'_o(P;t)$ for all $P \geq 0$, then KKT implies that it will never pay to use balancing power, i.e., $P_b^*(t) = 0$ at optimality. In this case, $P_d^*(t) = P_o^*(t) = \sum_i Q_i(t)$.

Hence welfare maximization reduces to the computation of the user consumptions $q_{ia}(t)$; the corresponding procurement decisions are then given by (22). The optimization of the social welfare in (15) then becomes:

$$\max_q \sum_{i,a} U_{ia}(q_{ia}) - \sum_t c \left(\sum_i Q_i(t); t \right) \quad (23)$$

$$\text{s. t. } A_{ia}q_{ia} \leq \eta_{ia}, \quad a \in \mathcal{A}_i, \forall i, \quad (24)$$

$$0 \leq Q_i(t) \leq \bar{Q}_i, \quad \forall i \quad (25)$$

The inequalities in (24) are the consumption constraints (1) of user i 's appliances and battery. The lower inequality in (25) is the same as (11); see the discussion at the end of Section 2.1 on battery constraints. The upper inequality in (25) imposes a bound on the total power drawn by user i . By assumption, the objective function is concave and the feasible set is convex. Hence an optimal point can in principle be computed offline centrally by the LSE. This however will require that the LSE know all the users' utility and battery cost functions and all the constraints, which is impractical for technical or privacy reasons. The goal of this section is to derive a distributed algorithm to solve (23)–(25) by decomposing it into subproblems that are solvable in a decentralized manner where the LSE only needs to know the aggregate demand but not the individual private information.

The key idea is for the LSE to set prices $\pi := (\pi(t), \forall t)$ to induce the users to individually choose socially optimal consumptions $q_i := (q_{ia}(t), \forall t)$ in response. Indeed, given prices π , we assume that each user i chooses its own demand q_i so as to maximize its net benefit, her total utility minus the electricity cost, i.e., each user i solves:

$$\max_{q_i} \sum_{a \in \mathcal{A}_i} U_{ia}(q_{ia}) - \sum_t \pi(t) Q_i(t) \quad \text{s. t. } (24) - (25). \quad (26)$$

Given prices π , we denote an *individually* optimal solution of (26) and the corresponding aggregate demand by

$$q_i(\pi) := (q_{ia}(t; \pi), \forall t, \forall a \in \mathcal{A}_i), \quad Q_i(\pi) := (Q_i(t; \pi), \forall t) := \left(\sum_{a \in \mathcal{A}_i} q_{i,a}(t; \pi), \forall t \right).$$

Recall $q(\pi) := (q_i(\pi), \forall i)$. It is a remarkable fact in the competitive equilibrium theory in economics that there exist prices π that align the individual optimality with the social optimality, i.e., there are prices π^* such that if $q_i(\pi^*)$ optimize i 's objectives for all users i then they also optimize the social welfare.

Definition 1. A consumption vector q^* is called (socially) *optimal* if it solves (23)–(25). A price vector π^* is called *optimal* if $q(\pi^*)$ is optimal, i.e., any solution $q(\pi^*)$ of (26) also solves (23)–(25).

The following result follows from the welfare theorem in economics. It implies that setting the prices to the marginal costs of power is optimal.

Theorem 1. *The prices that satisfy $\pi^*(t) := c'(\sum_i Q_i(t; \pi^*); t) \geq 0$ exist and are optimal.*

Proof. Write the welfare maximization problem as

$$\max_{q_i \in \mathcal{Q}_i, Y_i} \sum_{i,a} U_{ia}(q_{ia}) - \sum_t c \left(\sum_i Y_i(t); t \right) \quad \text{s. t.} \quad Y_i(t) = \sum_{a \in \mathcal{A}_i} q_{ia}(t), \quad \forall i, t$$

where the feasible set \mathcal{Q}_i is defined by the constraints (24)–(25). Clearly, an optimal solution q^* exists. Moreover, there exist Lagrange multipliers $\pi_i^*(t), \forall i, t$, such that (taking derivative with respect to $Y_i(t)$)

$$\pi_i^*(t) = c' \left(\sum_i Y_i^*(t); t \right) = c' \left(\sum_i \sum_{a \in \mathcal{A}_i} q_{ia}^*(t); t \right) \geq 0.$$

Since the right-hand side is independent of i , the LSE can set the prices as $\pi^*(t) := \pi_i^*(t) \geq 0$ for all i . One can check that the KKT condition for the welfare maximization problem are identical to the KKT conditions for the collection of users' problems. Since all these problems are convex, the KKT conditions are both necessary and sufficient for optimality. This proves the theorem. \square

3.2 Offline distributed scheduling algorithm

Theorem 1 motivates a distributed algorithm to compute the optimal prices π^* and user decisions $q(\pi^*)$. The LSE sets prices to be the marginal costs of power and each user solves its own maximization problem (26) in response. The model is that at the beginning of each day the LSE and (the energy management systems of) the users iteratively compute the electricity prices $\pi(t)$ and consumptions $q_i(t)$ for each period t of the following day. These decisions are then carried out for that day. This is an offline algorithm since all decisions are made at once before the day starts. It is decentralized where the LSE only knows the aggregate demand but not user utility functions or consumption constraints and the users do not need to coordinate among themselves but only respond to common prices.

Algorithm 1: Optimal scheduling without supply uncertainty

For each iteration $k = 1, 2, \dots$, after initialization:

1. The LSE collects aggregate demand forecasts, denoted by $(Q_i^k(t), \forall t)$, from all users i over a communication network. It updates the prices to the marginal costs $\pi^{k+1}(t) := c'(\sum_i Q_i^k(t); t)$ and broadcasts $\pi^{k+1} := (\pi^{k+1}(t), \forall t)$ to all users.
2. Each user i updates its demands q_i^{k+1} after receiving π^{k+1} according to

$$\begin{aligned} \tilde{q}_{ia}^{k+1}(t) &= q_{ia}^k(t) + \gamma \left(\frac{\partial U_{ia}(q_i^k)}{\partial q_{ia}^k(t)} - \pi^{k+1}(t) \right) \\ q_{ia}^{k+1} &= [\tilde{q}_{ia}^{k+1}]_{\mathcal{Q}_i} \end{aligned}$$

where $\gamma > 0$ is a constant stepsize, $\tilde{q}_{ia}^{k+1} := (\tilde{q}_{ia}^{k+1}(t), \forall t)$ is the new consumption vector before being projected onto the feasible set \mathcal{Q}_i specified by constraints (24)–(25), and $[\cdot]_{\mathcal{Q}_i}$ denotes this projection. User i 's aggregate demand forecast in period t is updated to $Q_i^{k+1}(t) = \sum_{a \in \mathcal{A}_i} q_{ia}^{k+1}(t)$.

3. Increment iteration index to $k + 1$ and goto Step 1.

Algorithm 1 converges asymptotically to optimal prices π^* and optimal consumptions $q(\pi^*)$, provided the stepsize $\gamma > 0$ is small enough. More precisely, suppose:

- A1: The utility functions $U_{ia}(q_{ia})$ are strictly concave in the vector $q_{ia} := (q_{ia}(t), \forall t)$ for all i, a .
- A2: The feasible set of q defined by the consumption constraints (24)–(25) is compact. All our user models in Section 2.1 satisfy this condition because of (6).
- A3: Suppose the spectral radius of the Hessian matrix $\nabla^2 U_{ia}$ and the second derivative $c''(\cdot; t)$ are both uniformly bounded: $\|\nabla^2 U_{ia}(q_{ia})\|_2 < \rho$ for all q_{ia} for all i, a , and $c''(Q; t) < \alpha$ for all Q, t .

Theorem 2. *Under the assumptions A1 – A3, the sequence (π^k, q^k) generated by Algorithm 1 converges to the optimal price and consumption vectors $(\pi^*, q(\pi^*))$, provided $\gamma < 2/(\rho + \alpha \sum_i |\mathcal{A}_i|)$.*

Proof. Let the welfare function be

$$h(q) := \sum_{i,a} U_{ia}(q_{ia}) - \sum_t c \left(\sum_i Q_i(t); t \right)$$

Then $h(q)$ is strictly concave since $U_{ia}(q_{ia})$ are strictly concave. The gradient $\nabla h(q)$ has components

$$[\nabla h(q)]_{ia}(t) = \frac{\partial U_{ia}(q_i)}{\partial q_{ia}(t)} - c' \left(\sum_i Q_i(t); t \right) \quad (27)$$

Hence Algorithm 1 is a gradient projection algorithm where in each iteration k , the variable q^k is updated to q^{k+1} according to:

$$q^{k+1} = \left[q^k + \gamma \nabla h(q^k) \right]_{\mathcal{Q}}$$

where $\mathcal{Q} := \mathcal{Q}_1 \times \cdots \times \mathcal{Q}_N$. Moreover assumption A3 implies the following lemma, proved in Appendix B.

Lemma 1. $\nabla h(q)$ is Lipschitz with $\|\nabla h(q) - \nabla h(\tilde{q})\|_2 < (\rho + \alpha \sum_i |\mathcal{A}_i|) \|q - \tilde{q}\|_2$ for all q, \tilde{q} .

Lemma 1 implies that, provided $\gamma < 2/(\rho + \alpha \sum_i |\mathcal{A}_i|)$, any accumulation point q^* of the sequence q^k generated by Algorithm 1 is optimal, i.e., maximizes welfare $h(q)$ [40, p. 214]. Assumption A2 implies that the sequence q^k lies in a compact set and hence must have a convergent subsequence. But assumption A1 implies that the optimal q^* is unique. Therefore all convergent subsequences, hence the original sequence q^k , must converge to q^* . By continuity of c' , $\pi^k(t) = c'(\sum_i Q_i^k(t); t)$ converges to the unique price $c'(\sum_i Q_i^*(t); t)$ with $Q_i^*(t) := \sum_{a \in \mathcal{A}_i} q_{ia}^*(t)$ which, by Theorem 1, is optimal. \square

The rate of convergence of Algorithm 1 depends on the stepsize γ : a larger γ generally leads to faster convergence, but a large γ can also risk instability. The bound on the stepsize γ in Theorem 2 is conservative; in practice a much larger stepsize can usually be used without losing stability. We simulate this algorithm in [39] with realistic system parameters. The simulation results show that, as expected, the prices are capable of coordinating the decisions of different appliances in a decentralized manner, to reduce peak aggregate demand and flatten its profile, greatly increasing the load factor. Furthermore, battery amplifies the benefits of demand response.

4 Conclusion

We have presented a simple yet versatile user model and formulated the optimal demand response problem as an $(1+T)$ -period dynamic program to maximize the expected social welfare. In this paper, we have focused on the case where there is no uncertainty. In this case demand response reduces to the deterministic welfare maximization in (23)–(25) that has a natural decentralized and incentive-compatible structure. We have proposed an offline distributed scheduling algorithm where the LSE sets the day-ahead prices to be their marginal costs based on forecast demands and, in response, the users forecast their demands to maximize their own surplus. As long as the stepsize is small enough, this procedure will converge to the unique optimal prices and consumptions. The algorithm is decentralized where the LSE only knows the aggregate demand but not user utility functions or consumption constraints, and the users do not need to coordinate with other users but only respond to the common prices from the LSE.

The current work has several limitations. First our model does not include the distribution system, implicitly assuming that the underlying network has enough capacity to distribute the power demanded by the users without causing congestion.

Second we only consider power balance in steady-state and ignore fast timescale dynamics such as frequency and voltage fluctuations due to random supply and demand. Third we do not model power market dynamics; for example, our model assumes that the cost functions faced by the LSE are independent of the demands and we ignore economic issues such as revenue-adequacy for the LSE. Finally our results are only for the case without uncertainty. When there is random renewable generation, offline scheduling alone will be insufficient and real-time demand response should be employed to match fluctuating supply. This is considered in [21].

Appendix A: Detailed appliance models

We describe detailed models of common electric appliances summarized in Section 2.1.

Type 1. This category of appliances includes lighting that must be on for a certain period of time. The consumption constraint is (6), with the understanding that $q_{ia}(t) = \bar{q}_{ia}(t) = 0$ for periods t that are outside its time of operation. User i attains a utility $U_{ia}(q_{ia}(t), t)$ from consuming power $q_{ia}(t)$ independent of its consumption in other periods, and the overall utility (2) is therefore separable in t .

Type 2. This category includes TV, video games, and computers. For these appliances, a user's utility depends on her consumption in each period she wishes to use it as well as the total amount of consumption in a day. Hence the consumption constraints are (6) and (7). For example, a user may have a favorite TV program that she wishes to watch everyday. With DVR, she can watch the program at any time. However the total power demand of TV should at least cover the program. Type 2 appliances have the same kind of utility functions (2) as Type 1 appliances. The time dependent utility function models the fact that a user may get different benefits from consuming the same amount of power at different times, e.g., she may enjoy a TV program to different levels at different times.

Type 3. This category includes PHEV, dish washer, clothes washer. For these appliances, a user only cares about whether the task is completed by a certain time. This means that the aggregate power consumption by such an appliance must exceed a threshold within its time of operation [28, 29, 33]. Hence the consumption constraints are (6) and (7). The utility depends only on the total power consumed, hence (3).

Type 4. This category includes HVAC (heating, ventilation, air conditioning) and refrigerator that control the temperature of a user's environment. Let $T_{ia}^{in}(t)$ and $T_{ia}^{out}(t)$ denote the temperatures at time t inside and outside the place that appliance (i, a) is in charge of, and \mathcal{T}_{ia} denotes the set of times when user i cares about the temperature. For instance, for air conditioner, $T_{ia}^{in}(t)$ is the temperature inside the house, $T_{ia}^{out}(t)$ is the temperature outside the house, and \mathcal{T}_{ia} is the set of times when she is at home.

The inside temperature evolves according to the following linear dynamics [27, 9, 26]:

$$T_{ia}^{in}(t) = T_{ia}^{in}(t-1) + \alpha(T_{ia}^{out}(t) - T_{ia}^{in}(t-1)) + \beta q_{ia}(t) \quad (28)$$

where α and β are parameters that specify thermal characteristics of the appliance and the environment in which it operates. The second term in equation (28) models heat transfer. The third term models the thermal efficiency of the system; $\beta > 0$ if appliance a is a heater and $\beta < 0$ if it is a cooler. Here, we define $T_{ia}^{in}(0)$ as the temperature $T_{ia}^{in}(T)$ from the previous day. Let $[\underline{T}_{ia}, \bar{T}_{ia}]$ be a range of preferred temperature, leading to the constraint:

$$\underline{T}_{ia} \leq T_{ia}^{in}(t) \leq \bar{T}_{ia}, \forall t \in \mathcal{T}_{ia}. \quad (29)$$

Using Equation (28), we can write $T_{ia}^{in}(t)$ in terms of $(q_{ia}(\tau), \tau = 1, \dots, t)$:

$$T_{ia}^{in}(t) = (1 - \alpha)^t T_{ia}^{in}(0) + \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} \alpha T_{ia}^{out}(\tau) + \beta \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} q_{ia}(\tau).$$

Define

$$T_{ia}(t) := (1 - \alpha)^t T_{ia}^{in}(0) + \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} \alpha T_{ia}^{out}(\tau). \quad (30)$$

Then

$$T_{ia}^{in}(t) = T_{ia}(t) + \beta \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} q_{ia}(\tau). \quad (31)$$

With (31), the constraint (29) becomes a linear constraint on the load vector q_{ia} : for any $t \in \mathcal{T}_{ia}$,

$$\underline{T}_{ia} \leq T_{ia}(t) + \beta \sum_{\tau=1}^t (1 - \alpha)^{t-\tau} q_{ia}(\tau) \leq \bar{T}_{ia}.$$

This is the constraint (8), in addition to (6). Assume user i attains a utility $U_{ia}(T_{ia}^{in}(t))$ when the temperature is $T_{ia}^{in}(t)$. Then (31) gives the utility function (4).

Appendix B: Proof of Lemma 1

We first describe the Hessian $\nabla^2 h(q)$. Let $N := |\mathcal{N}|$ be the number of users and $A := |\cup_{i \in \mathcal{N}} \mathcal{A}_i|$ the total number of appliances. Let k take value (i, a) for $i = 1, \dots, N, a = 1, \dots, A$. For $k = (i, a)$, let 1_k be 1 if $a \in \mathcal{A}_i$ and 0 otherwise. From (27), $\nabla^2 h(q)$ is given by

$$\begin{aligned}\frac{\partial^2 h}{\partial q_{\tilde{k}}^2(t)} &= \frac{\partial^2 U_k}{\partial q_{\tilde{k}}^2(t)}(q_k) - c'' \left(\sum_j \mathcal{Q}_j(t); t \right) \mathbf{1}_k \\ \frac{\partial^2 h}{\partial q_k(s) \partial q_k(t)} &= \frac{\partial^2 U_k}{\partial q_k(s) \partial q_k(t)}(q_k), \quad s \neq t \\ \frac{\partial^2 h}{\partial q_{\tilde{k}}(t) \partial q_k(t)} &= -c'' \left(\sum_j \mathcal{Q}_j(t); t \right) \mathbf{1}_k \mathbf{1}_{\tilde{k}}, \quad k \neq \tilde{k} \\ \frac{\partial^2 h}{\partial q_{\tilde{k}}(s) \partial q_k(t)} &= 0, \quad k \neq \tilde{k} \text{ and } s \neq t\end{aligned}$$

To express $\nabla^2 h(q)$ in matrix form, let $H_k(q_k)$ denote the $T \times T$ matrix $\frac{\partial^2 U_k}{\partial q_k^2}(q_k)$, for $k = 1, \dots, NA := K$. Let $H(q)$ denote the block-diagonal matrix

$$H(q) := \text{diag} (H_1(q_1), \dots, H_K(q_K))$$

Let C be the $TNA \times TNA$ matrix with $C_{kt, \tilde{k}\tilde{t}} := c'' \left(\sum_j \mathcal{Q}_j(t); t \right) \mathbf{1}_k \mathbf{1}_{\tilde{k}}$ if $t = \tilde{t}$ and 0 otherwise. Then $\nabla^2 h(q) = H(q) - C$. Hence $\|\nabla^2 h(q)\|_2 \leq \|H(q)\|_2 + \|C\|_2$.

Now assumption A3 implies

$$\|H(q)\|_2 \leq \max_k \|H_k(q_k)\|_2 \leq \rho$$

and (with $\tilde{k} = (\tilde{i}, \tilde{a})$)

$$\|C\|_2 = \rho(C) \leq \|C\|_\infty = \max_{kt} \sum_{\tilde{k}\tilde{t}} C_{kt, \tilde{k}\tilde{t}} \leq \alpha \max_k \mathbf{1}_k \sum_{\tilde{k}} \mathbf{1}_{\tilde{k}} = \alpha \sum_i |\mathcal{A}_i|$$

where $\rho(C)$ is the spectral radius of matrix C and the first equality holds because C is symmetric. Therefore $\|\nabla^2 h(q)\|_2 \leq \rho + \alpha \sum_i |\mathcal{A}_i|$. Theorem 9.19 of [41] implies that $\|\nabla h(q) - \nabla h(\tilde{q})\|_2 < (\rho + \alpha \sum_i |\mathcal{A}_i|) \|q - \tilde{q}\|_2$ for all q, \tilde{q} . \square

References

1. C. W. Gellings and J. H. Chamberlin. *Demand-Side Management: Concepts and Methods*. The Fairmont Press, 1988.
2. M. H. Albadi and E. F. El-Saadany. Demand response in electricity markets: An overview. In *Proceedings of the IEEE Power Engineering Society General Meeting*, June 2007.
3. A. I. Cohen and C. C. Wang. An optimization method for load management scheduling. *IEEE Transactions on Power Systems*, 3(2):612–618, May 1988.
4. Y. Y. Hsu and C. C. Su. Dispatch of direct load control using dynamic programming. *IEEE Transactions on Power Systems*, 6(3):1056–1061, August 1991.
5. D. C. Wei and N. Chen. Air conditioner direct load control by multi-pass dynamic programming. *IEEE Transactions on Power Systems*, 10(1):307–313, February 1995.
6. J. Chen, F. N. Lee, A. M. Breipohl, and R. Adapa. Scheduling direct load control to minimize system operation cost. *IEEE Transactions on Power Systems*, 10(4):1994–2001, November 1995.

7. K. H. Ng and G. B. Sheble. Direct load control – a profit-based load management using linear programming. *IEEE Transactions on Power Systems*, 13(2):688–695, May 1998.
8. W.-C. Chu, B.-K. Chen, and C.-K. Fu. Scheduling of direct load control to minimize load reduction for a utility suffering from generation shortage. *IEEE Transactions on Power Systems*, 8(4):1525–1530, November 1993.
9. B. Ramanathan and V. Vittal. A framework for evaluation of advanced direct load control with minimum disruption. *IEEE Transactions on Power Systems*, 23(4):1681–1688, November 2008.
10. M. D. Ilic, L. Xie, and J.-Y. Joo. Efficient coordination of wind power and price-responsive demand part I: Theoretical foundations; part II: Case studies. *IEEE Transactions on Power Systems*, 99, 2011.
11. Y. V. Makarov, C. Loutan, J. Ma, and P. de Mello. Operational impacts of wind generation on California power systems. *IEEE Transactions on Power Systems*, 24(2):1039–1050, May 2009.
12. M. C. Caramanis and J. M. Foster. Coupling of day ahead and real-time power markets for energy and reserves incorporating local distribution network costs and congestion. In *Proceedings of the 48th Annual Allerton Conference*, September – October 2010.
13. D. Kirschen. Demand-side view of electricity market. *IEEE Transactions on Power Systems*, 18(2):520–527, May 2003.
14. J. C. Smith, M. R. Milligan, E. A. DeMeo, and B. Parsons. Utility wind integration and operating impact: State of the art. *IEEE Transactions on Power Systems*, 22(3):900–908, August 2007.
15. N. Ruiz, I. Cobelo, and J. Oyarzabal. A direct load control model for virtual power plant management. *IEEE Transactions on Power Systems*, 24(2):959–966, May 2009.
16. P. P. Varaiya, F. F. Wu, and J. W. Bialek. Smart operation of smart grid: Risk-limiting dispatch. *Proceedings of the IEEE*, 99(1):40–57, January 2011.
17. Department of Energy. Benefits of demand response in electricity markets and recommendations for achieving them. Technical report, February 2006.
18. S. Borenstein. Time-varying retail electricity prices: Theory and practice. In Griffin and Puller, editors, *Electricity Deregulation: Choices and Challenges*. University of Chicago Press, 2005.
19. C. Triki and A. Violi. Dynamic pricing of electricity in retail markets. *Quarterly Journal of Operations Research*, 7(1):21–36, March 2009.
20. M. D. Ilic. Dynamic monitoring and decision systems for enabling sustainable energy services. *Proceedings of the IEEE*, 99(1):58–79, January 2011.
21. L. Jiang and S. H. Low. Optimal demand response: with uncertain supply. In *Technical Report*, 2011.
22. P. M. Schwarz, T. N. Taylor, M. Birmingham, and S. L. Dardan. Industrial response to electricity real-time prices: Short run and long run. *Economic Inquiry*, 40(4):597–610, 2002.
23. C. Goldman, N. Hopper, R. Bharvirkar, B. Neenan, R. Boisvert, P. Cappers, D. Pratt, and K. Butkins. Customer strategies for responding to day-ahead market hourly electricity pricing. Technical report, Lawrence Berkeley National Lab, LBNL-57128, August 2005. Report for CA Energy Commission.
24. T. N. Taylor, P. M. Schwarz, and J. E. Cochell. 24-7 hourly response to real-time pricing with up to eight summers of experience. *Journal of Regulatory Economics*, 27(3):235–262, January 2005.
25. S. Borenstein. The long-run efficiency of real-time electricity pricing. *The Energy Journal*, 26(3):93–116, February 2005.
26. J.E. Braun. Load control using building thermal mass. *Journal of solar energy engineering*, 125(3):292–301, August 2003.
27. P. Xu, P. Haves, M. A. Piette, and L. Zagreus. Demand shifting with thermal mass in large commercial buildings: Field tests, simulation and audits. Technical report, Lawrence Berkeley National Lab, LBNL-58815, 2006.
28. A. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE Transactions on Smart Grid*, 1(2):120–133, September 2010.

29. M. Pedrasa, T. Spooner, and I. MacGill. Coordinated scheduling of residential distributed energy resources to optimize smart home energy services. *IEEE Transactions on Smart Grid*, 1(2):134–143, September 2010.
30. S. Caron and G. Kesidis. Incentive-based energy consumption scheduling algorithms for the smart grid. In *Proceedings of the IEEE International Conference on Smart Grid Communications*, October 2010.
31. M. C. Caramanis and J. M. Foster. Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC)*, December 2009.
32. Z. Ma, D. Callaway, and I. Hiskens. Decentralized charging control for large populations of plug-in electric vehicles. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, December 2010.
33. K. Clement-Nyngs, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371–380, February 2010.
34. L. Chen, N. Li, and S.H. Low. Two Market Models for Demand Response in Power Networks. In *Proceedings of the IEEE International Conference on Smart Grid Communications*, October 2010.
35. G. Pritchard, G. Zakeri, and A. Philpott. A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations Research*, 58(4):1210–1219, July-August 2010.
36. M. J. Neely, A. S. Tehrani, and A. G. Dimakis. Efficient algorithms for renewable energy allocation to delay tolerant consumers. In *Proceedings of the IEEE International Conference on Smart Grid Communications*, October 2010.
37. Miao He, Sugumar Murugesan, and Junshan Zhang. Multiple timescale dispatch and scheduling for stochastic reliability in smart grids with wind generation integration. *preprint, CoRR*, abs/1008.3932, 2010.
38. C.-W. Tan and P. P. Varaiya. Interruptible electric power service contracts. *Journal of Economic Dynamics and Control*, 17(3):495–517, May 1993.
39. N. Li, L. Chen, and S. H. Low. Optimal Demand Response Based on Utility Maximization in Power Networks. In *Proceedings of IEEE Power Engineering Society General Meeting*, July 2011.
40. D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and distributed computation*. Old Tappan, NJ (USA); Prentice Hall Inc., 1989.
41. W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill Inc., 3 edition, 1976.