Simultaneous Allocation and Control of Distributed Energy Resources via Kullback-Leibler-Quadratic Optimal Control

Neil Cammardella[1], Ana Bušić[2], and Sean Meyn[3]

Abstract— **There is enormous flexibility potential in the power consumption of the majority of electric loads. This flexibility can be harnessed to obtain services for managing the grid: with carefully designed decision rules in place, power consumption for the population of loads can be ramped up and down, just like charging and discharging a battery, without any significant impact to consumers' needs. The concept is called *Demand Dispatch*, and the grid resource obtained from this design *virtual energy storage* (VES). In order to deploy VES, a balancing authority is faced with two challenges: 1. how to design local decision rules for each load given the target aggregate power consumption (distributed control problem), and 2. how to coordinate a portfolio of resources to maintain grid balance, given a forecast of net-load (resource allocation problem).**

Rather than separating resource allocation and distributed control, in this paper the two problems are solved simultaneously using a single convex program. The joint optimization model is cast as a finite-horizon optimal control problem in a mean-field setting, based on the new KLQ optimal control approach proposed recently by the authors.

The simplicity of the proposed control architecture is remarkable: With a large portfolio of heterogeneous flexible resources, including loads such as residential water heaters, commercial water heaters, irrigation, and utility-scale batteries, the control architecture leads to a single scalar control signal broadcast to every resource in the domain of the balancing authority.

Keywords: Smart grids, demand dispatch, distributed control, controlled Markov chains.

I. Introduction

A. Background

Aggressive renewable portfolio standards are set in place all over the world. Twenty nine states have adopted such standards within the U.S., with California and Hawaii targeting 100% renewable energy by 2045 [29]. The *global climate strike* of September 20, 2019 provides dramatic evidence that people all over the world will continue to push policy makers to incentivize clean energy.

While it is true that energy from sun and wind present challenges because of volatility and uncertainty of energy supply, we have faced similar challenges with balancing supply and demand in communication networks. Similar to managing supply and demand in the Internet and wireless communication, the solutions developed in this paper are based on distributed control, designed to maintain grid stability, while ensuring quality of service (QoS) to customers remains within specified bounds.

Policy makers today have emphasized battery technology to help absorb volatility — notably in California and Australia. Demand side resources can be used along side these battery systems, and with good design we will spend far less on these expensive resources [9], [10], [4].

Theory supporting distributed control of the power grid has developed rapidly over the past decade. The main idea is that many electric loads are highly flexible in terms of power consumption, and are already built with substantial computing power. This flexibility and "intelligence" can be harnessed to obtain services for managing the grid: with carefully designed decision rules in place, power consumption from the population of loads can be ramped up and down, just like charging and discharging a battery. The concept has been coined *Demand Dispatch* [3]. The grid resource obtained is called *virtual energy storage* (VES).

Consider for example an electric water heater – residential or commercial. In this case, maintaining QoS within desired limits simply means keeping the water temperature within given bounds. The load measures QoS directly; with the addition of a bit of communication, it can also receive a signal from a *balancing authority* or aggregator.

The balancing authority is then faced with two challenges: 1. how to design local decision rules for each load, and 2. how to coordinate a portfolio of resources to maintain grid balance, given a forecast of net-load (nominal load minus renewable generation).

The goal of [9] is to solve the second challenge. It is formulated as a quadratic program in which the objective function is a weighted sum of cost functions: some designed to smooth net-load to reduce stress on traditional generation, and others designed to ensure that loads in the population do not deviate significantly from nominal behavior. A surrogate for *state of charge* (SoC) for each class of loads is proposed, based on the virtual battery models of [23], [24], [20], [21], [13], [27]. This is a linear dynamical model, so the quadratic program is in fact a finite-horizon optimal control problem. Structure of the control solution was obtained in more recent research [26].

The output of that optimization problem is the optimal power deviation for each class of resources in the population. Five classes of flexible loads were considered in the numerical experiments of [9]: residential air conditioners (ACs), residential water heaters with faster cycle times (fHWs),

Funding provided through the National Science Foundation awards EPCN 1935389 & CPS 1646229, and French National Research Agency grant ANR-16-CE05-0008.

[1]Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611

[2]Inria and DI ENS, École Normale Supérieure, CNRS, PSL Research University, Paris, France

[3]Department of Electrical and Computer Engineering, University of Florida, and Inria International Chair, Paris

TABLE I
VES CAPACITY OF FIVE LOAD CLASSES IN CALIFORNIA

Par.	Unit	ACs	fWHs	sWHs	RFGs	PPs
N	million	7	0.7	0.7	17	1.2
C	GWh	5.6	0.26	0.70	0.43	14.55

commercial water heaters with slower cycle times (sHWs), refrigerators (RFGs), and pool pumps (PPs). The number of loads per class, N, and their flexible energy capacity, C, are listed in Table I; these values were chosen based on a residential energy consumption survey of California [33]. Aggregation is required to control the power consumption of 17 million refrigerators, and 7 million air conditioners.

In [9] it is argued that once this optimization problem is solved, it is up to the aggregator or balancing authority to decide how best to control the loads so that the aggregate power consumption in each class tracks the optimal target.

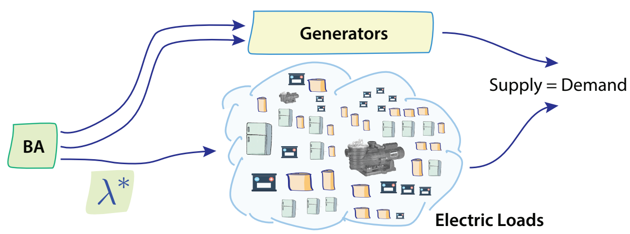

Fig. 1. Control architecture and KLQ optimal solution: a common scalar signal $\boldsymbol{\lambda}^*$ is broadcast to every resource in the population. This signal defines a local randomized decision rule at the resource.

In the present paper we adopt the framework of [27], [15], [6], [25] in which the aggregator broadcasts a common signal to each load in a given class. Communication from loads to aggregator is potentially useful, but only required for estimating the SoC [14]. However, instead of separating resource allocation and distributed control, in this paper the two problems are solved simultaneously using a single convex program. The optimization model is a generalization of the new approach to distributed control proposed in [8].

B. Contributions

The approach of this paper, in its simplest form, is the solution of a convex program with objective function

$$J(p) = \sum_{i=1}^{N_R} d(p^i, p^{0i}) + \frac{1}{2}\sum_{k=1}^{N_H}(y_k - r_k)^2 \tag{1}$$

where N_R denotes the number of resource classes and N_H is the terminal time; p^{0i} is a nominal probabilistic model for a load of class i, over the time horizon $[0, 1, \dots, N_H]$, p^i is the controlled model, and y_k is "virtual energy discharge" at time k based on these controlled models (see eq. (8)).

The "reference signal" $\boldsymbol{r} = \{r_k\}$ is much like the *balancing reserves* supplied by traditional generators in the Pacific Northwest [1]. However, the signals envisioned in the present paper are far more non-stationary and with far greater magnitude than anything seen in power grids today. Our goal is to eliminate ramps and peaks in net-load. So, for example, to balance California's grid today, the range of $\boldsymbol{r}$ may be nearly ± 10 GW.

In this paper, the "distance" $d(p^i, p^{0i})$ appearing in (1) will be defined using relative entropy (also known as Kullback-Leibler divergence). The objective (1) is a generalization of the KLQ cost criterion introduced in [8]. The optimizers $\{p^{*i}\}$ define a randomized decision rule for each load class.

It would appear that by performing the two optimization problems simultaneously – both resource allocation, and optimal distributed control – the overall computational burden would be insurmountable. This was our belief at the start.

Due to special structure of the KLQ cost criterion, and the fact that quadratic cost is only a function of the *sum* of power deviation, the optimization problem admits a Lagrangian decomposition, with special structure that is very attractive for implementation. The structure is illustrated in Fig. 1: The solution to the minimization of $J(p)$ is characterized by a single vector $\lambda^* \in \mathbb{R}^{N_H}$. The population may be extremely heterogeneous, including commercial air-conditioning, residential refrigerators, and even true storage devices such as batteries. Even though the probabilistic models are entirely different, each load performs an exponential "tilting" of its nominal behavior based on the common vector λ^*.

This paper introduces several variations of the basic control problem. In particular, it is found that it is valuable to introduce ramping costs, and transform techniques inspired by wavelets are used to reduce computational complexity.

C. Literature review

Since the seminal work of Schweppe [31] there has been great interest in decentralized control solutions to manage residential distributed energy resources, with much of this research focused on the use of price signals. Recent examples include [22], [34], [12] (and their extensive bibliographies).

Reference [16] is concerned with the 2012 ENTSO-E rules on demand-side management, that look similar to rules for "droop" for synchronous generators, and also Schweppe's proposal known as "FAPER" [31]. As expected based on prior research (see [30] for a bibliography), it is found that loads will sometime synchronize, leading to large oscillations in demand. This is why randomization is favored in both academic studies and recent patents such as [32].

Control techniques for demand-side management have evolved rapidly in the past decade. The *priority stack* approach [21] requires two way communication in real-time between aggregator and each load, but its simplicity makes it valuable when feasible. The approach of [2] uses two-way communication, but at a lower rate by using concepts from communication theory.

Decentralized control solutions through a control architecture similar to that proposed here are presented in [25]. Besides the difference in control architecture, the big difference between [16], [34] and the present paper is the choice of timescales. This work is about control at low frequencies, such as the balancing reserves in the U.S. Pacific Northwest, or the ramping/valley filling services at CAISO (see http://www.caiso.com for history since 2011).

Previous work based on an MPC architecture includes the aforementioned papers [4], [9], and [22], [19] which consider building operations, with attention to a range of QoS metrics. The load management problem is reduced to a linear program in [22] (eqns. (19,20)), and the goal is to obtain a decentralized solution in a game-theoretic setting. While [19], [11] aim to solve a game that arises in their control formulation, which is far from the motivation of the present work, the dual variables in these papers are similar to the signal $\boldsymbol{\lambda}^*$ that arises in this paper. Distributed techniques to learn $\boldsymbol{\lambda}^*$ may be devised based on the methods in this prior work, or consensus algorithms as in [18], [35], [5], [28].

The remainder of this paper is organized as follows: details on the optimization model for control is included in Section II, along with extensions. Section III contains a characterization of the solution to the general optimization problem, along with a proof – see Thm. 3.1. Numerical results are contained in Section IV, and conclusions and directions for further research in Section V.

II. Optimization Model

A. Notation

The control problem is formulated as a finite-horizon optimization problem over $[0, 1, \dots, N_H]$. In practice, this is just one part of an MPC control architecture.

- If μ is a pmf (probability mass function) on a discrete set S, and F a real valued function, its mean is denoted

$$\langle \mu, F\rangle = \sum_{x\in\mathsf{S}} \mu(x)F(x)$$

- N_R is the number of distributed energy resource (DER) classes; for each $i \in \{1, \dots, N_R\}$; there are n_i resources in DER class i.
- A load in class i is modeled as a stochastic process $\boldsymbol{X}^i$ evolving on the finite state space denoted $\mathsf{X}^i = \{x^{i1}, \dots, x^{id}\}$. S^i denotes the set of sequences in X^i:

$$\mathsf{S}^i = \mathsf{X}^i \times \cdots \times \mathsf{X}^i \qquad (N_H + 1 \text{ times}).$$

The common distribution is denoted: $p^i(\vec{x}) := \mathsf{P}\{\boldsymbol{X}^i = \vec{x}\}$, $\vec{x} \in \mathsf{S}^i$. That is,

$$p^i(x_0, x_1, \dots, x_{N_H}\} = \mathsf{P}\{X_0^i = x_0, \dots X_{N_H}^i = x_{N_H}\}$$

The kth marginal is denoted $\nu_k^i(x) := \mathsf{P}\{X_k^i = x\}$, $x \in \mathsf{X}^i$.

- If $\boldsymbol{X}^i$ is Markovian, then we let $\{P_k^i : k \ge 0\}$ denote the transition matrices. In this case, the marginals evolve according to the linear "state equations":

$$\nu_{k+1}^i = \nu_k^i P_k^i \qquad k \ge 0\,. \tag{2}$$

- The *nominal model* describes behavior without interaction with the balancing authority, and is assumed to be Markovian. This is distinguished using the notation $\{\nu_k^{0i}\,,\ P_k^{0i} : k \ge 0\}$, and p^{0i}. The Markov Property implies the factorization:

$$p^{0i}(\vec{x}) = \nu_0^{0i}(x_0^i)P_0^{0i}(x_0^i, x_1^i)P_1^{0i}(x_1^i, x_2^i)\cdots \tag{3}$$

The overall stochastic models are denoted $p^0 := \{p^{0i} : 1 \le i \le N_R\}$ and $p := \{p^i : 1 \le i \le N_R\}$.

- For each $1 \le i \le N_R$, power consumption as a function of state is denoted $\mathcal{U}^i \colon \mathsf{X}^i \to \mathbb{R}_+$, and its mean under the nominal model

$$\bar{\mathcal{U}}_k^{0i} = \langle \nu_k^0, \mathcal{U}^i\rangle$$

The "virtual energy discharge" at state x and time k, is

$$\mathcal{Y}_k^i(x) = -\Delta[\mathcal{U}^i(x) - \bar{\mathcal{U}}_k^{0i}] \tag{4}$$

where Δ is the inter-sampling time.

- Kullback-Leibler divergence, or relative entropy, is the mean log-likelihood:

$$D(p^i \| p^{0i}) = \sum_{\vec{x}\in\mathsf{S}^i} \log\Big(\frac{p^i(\vec{x})}{p^{0i}(\vec{x})}\Big)\, p^i(\vec{x}) \tag{5}$$

It is assumed in this paper that the nominal model p^0 is obtained through a combination of system identification and *design*. The design aspect is based on the construction of a nominal decision rule for the load. The special case of an air conditioner is illustrated in Fig. 2. The standard hysteresis band $[\Theta^{\min}, \Theta^{\max}]$ is replaced by *randomized* decision rule: if the AC is on, and the room temperature is θ degrees at time slot k, then the load is turned off with probability $p^{\ominus}(\theta)$.

If the state space for an AC is chosen to be $\mathsf{X}^i = \{0, 1\} \times \mathbb{R}$, where the discrete variable is power mode and the continuous variable is temperature, then these switching probabilities define a randomized policy:

$$\phi^{0i}(0 \mid x) = \begin{cases} p^{\ominus}(\theta) & x = (1, \theta) \\ 1 - p^{\oplus}(\theta) & x = (0, \theta) \end{cases}$$

and $\phi^{0i}(1 \mid x) = 1 - \phi^{0i}(0 \mid x)$.

A full construction of the nominal model is described as follows. Suppose we have deterministic dynamics for a nominal model without randomization:

$$(U_{t_{k+1}}, \theta_{t_{k+1}}) = f_k^i(U_{t_k}, \theta_{t_k})$$

where U denotes power mode, and θ temperature. In this case, denoting $x = (m, \theta)$ and $x' = (m', \theta')$, we take

$$\begin{aligned} P_k^{0i}(x, x') = \phi^{0i}(0 \mid x)\mathbb{I}\{\theta' = f_k^i(0, \theta)\} \\ + \phi^{0i}(1 \mid x)\mathbb{I}\{\theta' = f_k^i(1, \theta)\} \end{aligned} \tag{6}$$

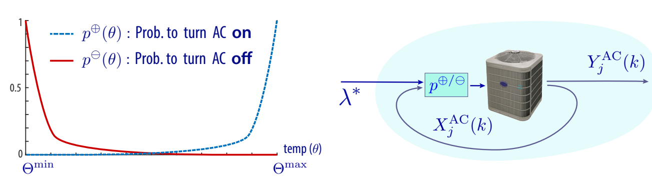

Fig. 2. Switching probabilities and local decision making for an AC. The two plots show the *nominal* probability of switching power mode as a function of state. The signal $\boldsymbol{\lambda}^*$ will modify this nominal behavior. While the sequence $\boldsymbol{\lambda}^*$ is identical for each load in the population, the transformation is sensitive to the nominal load dynamics.

It is assumed that the only randomness in the nominal model is through the randomized decision rule, so in particular the construction (6) is used for TCLs. This ensures that the optimal control solution is feasible. The fully probabilistic solution is the subject of current research [7].

B. Control Objective

The objective function is expressed as follows:

$$J(p;\nu_0^0) := \sum_{i=1}^{N_R} D(p^i\|p^{0i}) + \frac{1}{2}\sum_{n=1}^{N_\gamma} \kappa_n\gamma_n^2 \tag{7}$$

in which $D(p^i\|p^{0i})$ denotes Kullback-Leibler divergence (interpreted as the *control cost*), and the quadratic terms are designed to encourage reference tracking. The integer N_γ and non-negative weights $\{\kappa_n\}$ are design parameters, and each γ_n is an affine function of p. Examples are provided in the following.

Basic model: For given p, the *virtual energy discharge* is denoted (recalling (4)),

$$y_k = \sum_{i=1}^{N_R} n_i \sum_{x\in\mathsf{X}^i} \nu_k^i(x)\mathcal{Y}_k^i(x)\,, \quad 0\le k\le N_H \tag{8}$$

If $y_k>0$, this means that the loads are consuming less power at time k as compared to consumption under the nominal model p^0. The constraint $\nu_0=\nu_0^{0i}$ is imposed throughout the paper, which implies that $y_0=0$ for any feasible p.

In the basic model we take $N_\gamma=2N_H$, and

$$\begin{aligned}\gamma_k &= y_k - r_k\\ \gamma_{N_H+k} &= y_k - y_{k-1} - (r_k - r_{k-1})\,, \quad 1\le k\le N_H\end{aligned} \tag{9}$$

The objective function (7) is then a weighted combination of relative entropy, cost on deviation, and cost for ramping. The objective function (1) is a special case of the basic model in which $\kappa_n=1$ for $n\le N_H$, and $\kappa_n=0$ for $n>N_H$.

The difficulty with this formulation is complexity: for a 24-hour time horizon, with 5-minute sampling, this results in $N_H=576$. The resulting optimization problem is ill-conditioned, but may be easily solved with a carefully designed algorithm. However, there is little motivation for exactly solving the basic model since there are attractive alternatives.

Transform techniques: Lossy compression of $\boldsymbol{r}$ can be used to reduce complexity. The transform techniques are based on a collection of functions $\{w_n : 1\le n\le N_W\}$, with $w_n\colon\{0,1,\dots,N_H\}\to\mathbb{R}$ for each n, and $N_W\ll N_H$. It is assumed that these functions are orthonormal:

$$\sum_k w_n(k)w_{n'}(k) = \mathbb{I}\{n=n'\} \tag{10}$$

Standard examples are Fourier series and also the *degenerate* case

$$w_n^\bullet(k) := \mathbb{I}\{n=k\}\,, \qquad 1\le n,k\le N_H \tag{11}$$

The transformed signal is the N_W-dimensional vector $\hat{\boldsymbol{r}}$, with

$$\hat r_n = \sum_{k=1}^{N_H} w_n(k) r_k\,, \qquad 1\le n\le N_W\,. \tag{12}$$

The functions on X^i used in the definition of $\{y_k\}$ are similarly transformed functions defined on S^i. Denote

$$\widehat{\mathcal{Y}}_n^i(\vec{x}) = \sum_{k=1}^{N_H} w_n(k)\mathcal{Y}^i(x_k^i) \tag{13a}$$

$$^{\delta}\widehat{\mathcal{Y}}_n^i(\vec{x}) = \sum_{k=1}^{N_H} w_n(k)[\mathcal{Y}^i(x_k^i) - \mathcal{Y}^i(x_{k-1}^i)] \tag{13b}$$

and for any model p denote

$$\hat y_n = \sum_{i=1}^{N_R} n_i\langle p^i, \widehat{\mathcal{Y}}_n^i\rangle \tag{14a}$$

$$^{\delta}y_n = \sum_{i=1}^{N_R} n_i\langle p^i, {}^{\delta}\widehat{\mathcal{Y}}_n^i\rangle\,, \quad 0\le n\le N_W \tag{14b}$$

An application of the transforms results in $N_\gamma=2N_W$ for the optimization model (7), and for each $1\le n\le N_W$,

$$\gamma_n = \hat y_n - \hat r_n \tag{15a}$$

$$\gamma_{N_W+n} = {}^{\delta}y_n - (\hat r_n - \hat r_{n-1}) \tag{15b}$$

III. KLQ and Lagrangian Decomposition

In analysis and code it is convenient to restrict only to γ_n of the form (15a) through the following transformations: define for $1\le n\le N_W$, and $1\le k\le N_H$,

$$\begin{aligned}\hat r_{N_W+n} &= \hat r_n - \hat r_{n-1}\\ w_{N_W+n}(k) &= -w_{N_W+n}(k+1) + w_{N_W+n}(k)\end{aligned} \tag{16}$$

with the convention that $w_n(N_W+1)=0$. We have thus doubled the number of transform functions. With Eqs. (13a), (14a) and (15a) defined with respect to the extended set of $\{w_n\}$, we obtain

$$\gamma_{N_W+n} := {}^{\delta}y_n - (\hat r_n - \hat r_{n-1}) = \hat y_{N_W+n} - \hat r_{N_W+n}$$

In the remaining analysis we assume that all γ_n are of the form (15a) for some w_n, and relax the requirement (10). In this case, with the objective function expressed

$$J(p;\nu_0^0) \ := \sum_{i=1}^{N_R} D(p^i\|p^{0i}) + \frac{1}{2}\sum_{n=1}^{N_\gamma}\kappa_n\gamma_n^2\,, \tag{17}$$

the goal is to minimize over p subject to

$$\begin{aligned}\gamma_n &= \hat y_n - \hat r_n\,, \quad 1\le n\le N_\gamma\\ \nu_0^i &= \nu_0^{0i}\,, \quad 1\le i\le N_R.\end{aligned} \tag{18}$$

The minimum of (17) over all pmfs p satisfying (18) is denoted $J^*(\nu_0^0)$.

The Lagrange multipliers for the first and second constraints are denoted $\lambda\in\mathbb{R}^{N_\gamma}$ and $h\in\mathbb{R}^{N_R}\times\mathbb{R}^d$, respectively. Although h^i, the ith element of h, can be regarded as a real-valued function on X^i, we take $h^i\colon\mathsf{S}^i\to\mathbb{R}$ to

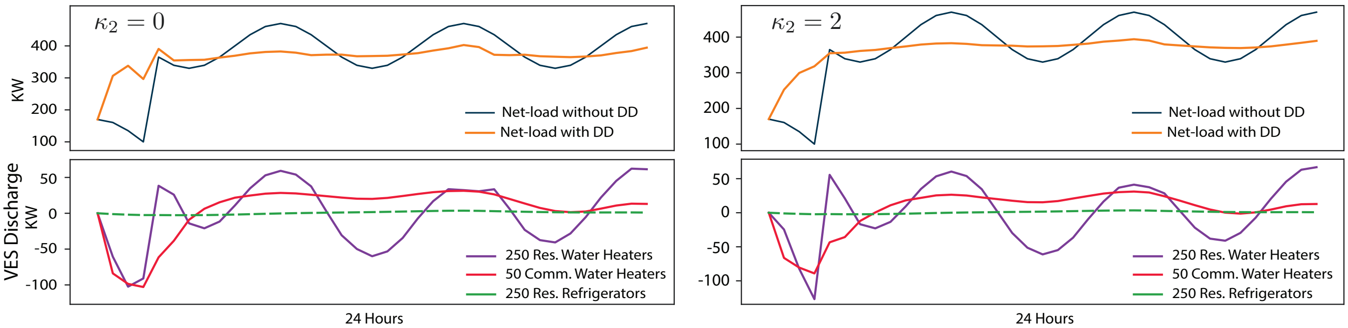

Fig. 3. Balancing with VES: impact of the ramping penalty κ_2

simplify exposition; that is, $h^i(x_0, x_1, x_2, \dots)$ depends only on x_0. The Lagrangian is the augmented function:

$$
\begin{aligned}
\mathcal{L}(p,\gamma,\lambda,h) = &\sum_{i=1}^{N_R} D(p^i \| p^{0i}) + \frac{1}{2}\sum_{n=1}^{N_\gamma} \kappa_n \gamma_n^2 \\
&+ \sum_{i=1}^{N_R} \langle p^i - p^{0i}, h^i \rangle \\
&+ \sum_{n=1}^{N_\gamma} \lambda_n [\gamma_n + \hat{r}_n - \hat{y}_n]
\end{aligned}
\tag{19}
$$

and its infimum over p and γ is the dual function:

$$
\varphi^*(\lambda, h) = \inf_{p,\gamma} \mathcal{L}(p,\gamma,\lambda,h)\ . \tag{20}
$$

The supremum over h is denoted

$$
\varphi^*(\lambda) := \sup_h \varphi^*(\lambda, h) \tag{21}
$$

Slater's condition implies that strong duality holds so $\max_{\lambda,h} \varphi^*(\lambda,h) = J^*(\nu_0^0)$.

For any $\lambda \in \mathbb{R}^{N_\gamma}$ and any i, let $p^{\lambda i}$ denote the pmf with log-likelihood ratio $L^{\lambda i} = \log(p^{\lambda i}/p^{0i})$ given by

$$
L^{\lambda i}(\vec{x}) = \sum_{n=1}^{N_\gamma} \lambda_n n_i \widehat{\mathcal{Y}}_n^i(\vec{x}) - \Lambda_\lambda^i(x_0)\,, \tag{22}
$$

where for each $x_0 \in \mathsf{X}^i$,

$$
\Lambda_\lambda^i(x_0) = \log\Bigl[\sum_{\vec{x}} p^{0i}(\vec{x} \mid x_0) \exp\Bigl\{\sum_{n=1}^{N_\gamma} \lambda_n n_i \widehat{\mathcal{Y}}_n^i(\vec{x})\Bigr\}\Bigr] \tag{23}
$$

Observe that $p^{\lambda i}$ and p^{0i} have common first marginals.

Theorem 3.1. *Consider the problem of minimizing $J(p;\nu_0^0)$ over the set of pmfs p with given initial marginals ν_0^0. Assume moreover that $\kappa_n > 0$ for each n. Then:*

(i) *An optimizer p^* exists and is unique.*

(ii) *The log-likelihood ratios $L^{*i} = \log(p^{*i}/p^{0i})$ are given by*

$$
L^{*i}(\vec{x}) = \sum_{n=1}^{N_\gamma} \lambda_n^* n_i \widehat{\mathcal{Y}}_n^i(\vec{x}) - \Lambda_{\lambda^*}^i(x_0^i) \tag{24}
$$

(iii) *The function defined in (21) is strictly concave, and admits the representation*

$$
\varphi^*(\lambda) = -\frac{1}{2}\sum_{n=1}^{N_\gamma} \kappa_n \lambda_n^2 + \lambda^T \hat{r} - \sum_{i=1}^{N_R} \langle \nu^{0i}, \Lambda_\lambda^i \rangle \tag{25}
$$

with partial derivatives, for $0 \le n \le N_\gamma$,

$$
\frac{\partial}{\partial \lambda_n} \varphi^*(\lambda) = -\frac{1}{\kappa_n}\lambda_n + \hat{r}_n - \sum_{i=1}^{N_R} n_i \langle p^{\lambda i}, \widehat{\mathcal{Y}}_n^i \rangle \tag{26}
$$

where $p^{\lambda i}$ has log likelihood ratio (22).

(iv) *The optimizer λ^* satisfies, for each $0 \le n \le N_\gamma$,*

$$
0 = -\frac{1}{\kappa_n}\lambda_n^* + \hat{r}_n - \sum_{i=1}^{N_R} n_i \langle p^{*i}, \widehat{\mathcal{Y}}_n^i \rangle \tag{27}
$$

Proof. To establish (i) we show that the Lagrangian admits a unique optimizer $p^{\lambda,h}$ for each λ and h. The assertion (i) follows once we obtain a unique optimizer (λ^*, h^*) for the dual function.

Express the Lagrangian as

$$
\begin{aligned}
\mathcal{L}(p,\gamma,\lambda,h) = &\sum_{i=1}^{N_R} \mathcal{L}^i(p^i,\lambda,h^i) - \sum_{i=1}^{N_R} \langle p^{0i}, h^i \rangle \\
&+ \lambda^T(\gamma + \hat{r}) + \frac{1}{2}\sum_{n=1}^{N_\gamma} \kappa_n \gamma_n^2\,,
\end{aligned}
$$

with $\quad \mathcal{L}^i(p^i,\lambda,h^i) = D(p^i \| p^{0i}) + \langle p^i, h^i - n_i \lambda_n \widehat{\mathcal{Y}}_n^i \rangle$

We thus obtain a Lagrangian decomposition:

$$
\begin{aligned}
\mathcal{L}(p,\gamma,\lambda,h) = &\sum_{i=1}^{N_R} \inf_{p^i} \mathcal{L}^i(p^i,\lambda,h^i) - \sum_{i=1}^{N_R} \langle p^{0i}, h^i \rangle \\
&+ \inf_\gamma \Bigl[\lambda^T(\gamma + \hat{r}) + \frac{1}{2}\sum_{n=1}^{N_\gamma} \kappa_n \gamma_n^2\Bigr]
\end{aligned}
\tag{28}
$$

The minimization over γ is trivial: $\gamma_n^* = \kappa_n^{-1}\lambda_n$. The minimization over p^i appears more complex, but this can be expressed in the suggestive form:

$$
\inf_{p^i} \mathcal{L}^i(p^i,\lambda,h^i) = -\sup_{p^i} \Bigl[\langle p^i, F^i \rangle - D(p^i \| p^{0i})\Bigr], \tag{29}
$$

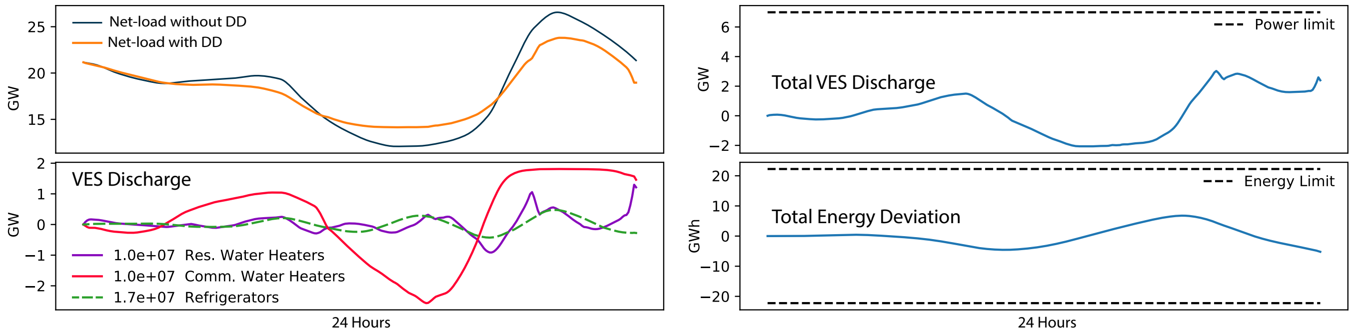

Fig. 4. Balancing the California "Duck Curve". Peaks and ramps are significantly reduced without violating customers' QoS.

with $F^i = n_i \sum_n \lambda_n \widehat{\mathcal{Y}}_n^i - h^i$. It follows from Kullback's Lemma [17] that there is a unique optimizer

$$p^{\lambda,h,i} = p^{0i}(\vec{x}) \exp\left(F^i(\vec{x}) - \Lambda(F^i)\right) \tag{30}$$

where $\Lambda(F^i)$ is a normalizing constant.

For a given λ, maximizing $\varphi^*(\lambda, h)$ over h^i leads to

$$h_\lambda^i(\vec{x}) = \Lambda_\lambda^i(x_0) + \text{Const.},$$

where the constant is arbitrary (a constant does not change the value of the Lagrangian in (19)). The identity (25) is obtained on substituting $p^{\lambda,h,i}$, $\gamma_n^* = \kappa_n^{-1}\lambda_n$ and $h_\lambda^i = \Lambda_\lambda^i$ into (28).

The function Λ_λ^i appearing in (23) is convex in λ (for each x_0) since it is a log-moment generating function. Under the assumption that $\kappa_n > 0$ for each n, it follows that the function $\varphi^*(\lambda)$ in (25) is strictly concave, and admits a unique maximizer λ^* that satisfies the first-order optimality conditions (27) for each n.

Part (i) follows as claimed, with $p^{*i} = p^{\lambda^*,h^*,i}$. □

IV. Numerical Experiments

The dual function (25) is strictly convex and its gradient (26) is easily computed. In the following numerical experiments, proximal gradient methods were used to obtain λ^* and h^*, and hence, p^*. The reference signal was chosen to be net load and it was centered by subtracting its mean. The goal of these experiments was to learn what kind of trajectories the loads can easily provide, and to find the resulting net-load that must be supplied by traditional generation. The difference in the net load curves with and without DD is due to the VES "discharge", which is simply the negative of the collective power deviation for each class. Three classes of homogeneous electric loads were included: residential water heaters, commercial water heaters, and residential refrigerators.

The first experiment, displayed in Fig. 3, was chosen to illustrate the potential for flexible loads to 'flatten' a net load curve and to demonstrate the benefits of including a ramping penalty in the objective function. In both cases, the net load was significantly flatter. However, with the introduction of a ramping penalty, the large ramp in the net load was eliminated.

A second experiment was performed using a more realistic net load profile. The 'Duck Curve', shown in Fig. 4, is a typical net load profile that would be experienced in California on a sunny day. During the middle of the day, the net load is very low due to an abundance of solar generation, but in the evening, as the sun sets and people turn on their appliances, the net load can ramp up as fast as 10 GW in three hours! This presents a major challenge for grid operators, many of which are already spending millions of dollars on lithium-ion batteries to "flatten the Duck". However, as shown in Fig. 4, demand dispatch was able to reduce the peak by more than two GW and significantly reduced the evening ramp. Notice that the VES "discharge" of the collection was well within its power and energy limits [21]; the deviation from nominal behavior of the electric loads would be nearly invisible to the customers since QoS would be satisfied for every individual.

V. Conclusions

The work presented in this paper shows that distributed energy resources can be optimally allocated and controlled using a single convex program. Surprisingly, every class of DERs receives the same control signal, and although their responses can very wildly, their cumulative effect is a flatter, smoother net load curve.

As observed in prior research, it is found that the controlled behavior of the aggregate of electric loads appears very similar to a grid-scale battery, but are potentially much cheaper. The main contribution here is the new methodology to simultaneously perform resource allocation and control.

Many extensions are planned for future work. For example, this control scheme assumes each class of DERs is homogeneous; what happens if this is only an approximation? Robustness to heterogeneity, modeling error, and performance over long time scales can be improved with the introduction of feedback, such as through model predictive control (MPC). A combination of non-uniform sampling and alternative transform techniques are being investigated as ways to create effective and computationally efficient MPC implementations.

References

[1] BPA balancing authority. Online, http://tinyurl.com/BPAgenload, http://tinyurl.com/BPAbalancing.

[2] M. Almassalkhi, J. Frolik, and P. Hines. Packetized energy management: asynchronous and anonymous coordination of thermostatically controlled loads. In *Proc. of the American Control Conf.*, pages 1431–1437. IEEE, 2017.

[3] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl. Demand dispatch. *IEEE Power and Energy Magazine*, 8(3):20–29, May 2010.

[4] J. Brooks and P. Barooah. Virtual energy storage through decentralized load control with quality of service bounds. In *Proc. of the American Control Conf.*, pages 735–740, May 2017.

[5] J. Brooks and P. Barooah. Consumer-aware distributed demand-side contingency service in the power grid. *IEEE Transactions on Control of Network Systems*, 5(4):1987–1997, Dec 2018.

[6] A. Bušić and S. Meyn. Distributed randomized control for demand dispatch. In *Proc. of the IEEE Conf. on Dec. and Control*, pages 6964–6971, Dec 2016.

[7] N. Cammardella, A. Bušić, and S. Meyn. Kullback-Leibler-quadratic optimal control. *ArXiv e-prints*, page PP, March 2020.

[8] N. Cammardella, A. Bušić, Y. Ji, and S. Meyn. Kullback-Leibler-Quadratic optimal control of flexible power demand. In *Proc. of the IEEE Conf. on Dec. and Control*, page PP, Dec. 2019.

[9] N. Cammardella, J. Mathias, M. Kiener, A. Bušić, and S. Meyn. Balancing California's grid without batteries. In *Proc. of the IEEE Conf. on Dec. and Control*, pages 7314–7321, Dec 2018.

[10] N. Cammardella, B. Moye, Y. Chen, and S. Meyn. An energy storage cost comparison: Li-ion batteries vs distributed load control. In *Clemson University Power System Conference*, pages 1–6, Sep 2018.

[11] R. Carli and M. Dotoli. A decentralized resource allocation approach for sharing renewable energy among interconnected smart homes. In *IEEE Conference on Decision and Control*, pages 5903–5908, Dec 2015.

[12] P. Chakraborty, E. Baeyens, and P. P. Khargonekar. Grid integration of renewable electricity and distributed control. In R. Tempo, S. Yurkovich, and P. Misra, editors, *Emerging Applications of Control and Systems Theory: A Festschrift in Honor of Mathukumalli Vidyasagar*, pages 205–216. Springer International Publishing, Cham, 2018.

[13] Y. Chen, A. Bušić, and S. Meyn. Individual risk in mean field control with application to automated demand response. In *Proc. of the IEEE Conf. on Dec. and Control*, pages 6425–6432, Dec 2014.

[14] Y. Chen, A. Bušić, and S. Meyn. State estimation for the individual and the population in mean field control with application to demand dispatch. *IEEE Transactions on Automatic Control*, 62(3):1138–1149, March 2017.

[15] Y. Chen, M. U. Hashmi, J. Mathias, A. Bušić, and S. Meyn. Distributed control design for balancing the grid using flexible loads. In S. Meyn, T. Samad, I. Hiskens, and J. Stoustrup, editors, *Energy Markets and Responsive Grids: Modeling, Control, and Optimization*, pages 383–411. Springer, New York, NY, 2018.

[16] F. Conte, S. Massucco, F. Silvestro, E. Ciapessoni, and D. Cirio. Stochastic modelling of aggregated thermal loads for impact analysis of demand side frequency regulation in the case of Sardinia in 2020. *International J. of Electrical Power & Energy Systems*, 93:291 – 307, 2017.

[17] A. Dembo and O. Zeitouni. *Large Deviations Techniques And Applications*. Springer-Verlag, New York, second edition, 1998.

[21] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Aggregate flexibility of thermostatically controlled loads. *IEEE Trans. on Power Systems*, 30(1):189–198, Jan 2015.

[18] A. D. Domínguez-García and C. N. Hadjicostis. Distributed algorithms for control of demand response and distributed energy resources. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 27–32, Dec 2011.

[19] H. Hao, J. Lian, K. Kalsi, and J. Stoustrup. Distributed flexibility characterization and resource allocation for multi-zone commercial buildings in the smart grid. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 3161–3168, Dec 2015.

[20] H. Hao, B. Sanandaji, K. Poolla, and T. Vincent. A generalized battery model of a collection of thermostatically controlled loads for providing ancillary service. In *51st Annual Allerton Conference on Communication, Control, and Computing*, pages 551–558, Oct 2013.

[22] L. A. Hurtado, E. Mocanu, P. H. Nguyen, M. Gibescu, and W. L. Kling. Comfort-constrained demand flexibility management for building aggregations using a decentralized approach. In *2015 International Conference on Smart Cities and Green ICT Systems (SMARTGREENS)*, pages 1–10, May 2015.

[23] K. Kalsi, F. Chassin, and D. Chassin. Aggregated modeling of thermostatic loads in demand response: A systems and control perspective. In *Proc. of the 50th IEEE Conf. on Dec. and Control*, pages 15–20, 2011.

[24] K. Kalsi, M. Elizondo, J. Fuller, S. Lu, and D. Chassin. Development and validation of aggregated models for thermostatic controlled loads with demand response. In *45th Hawaii International Conference on System Sciences (HICSS)*, volume 18, pages 1959–1966, Jan. 2012.

[25] J. Mathias, A. Bušić, and S. Meyn. Demand dispatch with heterogeneous intelligent loads. In *Proc. 50th Annual Hawaii International Conference on System Sciences (HICSS), and arXiv 1610.00813*, 2017.

[26] J. Mathias, R. Moye, S. Meyn, and J. Warrington. State space collapse in resource allocation for demand dispatch. In *Proc. of the IEEE Conf. on Dec. and Control*, page PP, Dec. 2019.

[27] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren. Ancillary service to the grid using intelligent deferrable loads. *IEEE Trans. Automat. Control*, 60(11):2847–2862, Nov 2015.

[28] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei. A survey of distributed optimization and control algorithms for electric power systems. *IEEE Transactions on Smart Grid*, 8(6):2941–2962, Nov 2017.

[29] National Conference of State Legislatures. State renewable portfolio standards and goals. http://tinyurl.com/nhyhl64. Accessed: 2019-09-12.

[30] M. S. Nazir and I. A. Hiskens. Load synchronization and sustained oscillations induced by transactive control. In *IEEE Power Energy Society General Meeting*, pages 1–5, July 2017.

[31] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox. Homeostatic utility control. *IEEE Trans. on Power Apparatus and Systems*, PAS-99(3):1151 –1163, May 1980.

[32] J. Sharp. Electrical load disconnect device with electronic control, Dec. 2012. US Patent 8,328,110.

[33] United States Energy Information Administration. *2009 Residential Energy Consumption Survey*.

[34] S. Williams, M. Short, and T. Crosbie. On the use of thermal inertia in building stock to leverage decentralised demand side frequency regulation services. *Applied Thermal Engineering*, 133:97 – 106, 2018.

[35] C. Zhao, U. Topcu, and S. H. Low. Optimal Load Control via Frequency Measurement and Neighborhood Area Communication. *IEEE Trans. Power Systems*, 28:3576 – 3587, November 2013.