This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2023.3293037

State Space Collapse in Resource Allocation for Demand Dispatch and its Implications for Distributed Control Design

Joel Mathias, Sean Meyn, Robert Moye, and Joseph Warrington

***Abstract*—It is now well established that many electric loads are inherently flexible, and this flexibility can be harnessed to provide grid services identical to those obtained today through batteries and responsive generators. This paper concerns the resource allocation problem associated with control of a large collection of heterogeneous loads. This problem is posed as a finite-horizon optimal control problem, in which the cost function reflects both the needs of the grid and the needs of the users of electric loads in the population. The main result is a form of *state space collapse*: the marginal cost for each load class evolves in a two-dimensional subspace, spanned by a scalar co-state process and its derivative.**

I. Introduction

The Federal Energy Regulatory Commission (FERC) Order No. 2222, approved in September 2020, requires the development of rules and incentives to integrate distributed energy resources (DERs) in grid operations [15]. This development along with policy promoted by federal and state elected officials bring concomitant challenges associated with the volatility and uncertainty of DERs and renewable generation.

Volatility is evident in Fig. 1, showing data from CAISO on Mar 16, 2021: solar and wind supply nearly 15 GW of power during the afternoon, while they supply less than 3 GW at night. Renewable energy briefly served 90% of the electricity demand in the Southwest Power Pool grid during the early morning hours of March 29, 2022 [16]; market operators indicated that over 88% of total demand was served by wind at some point in time. Addressing these peaks and ramps through traditional generation (e.g. coal and natural gas) is costly [25].

These challenges are one reason for interest in techniques to better manage the grid through new battery technology, as well as the application of tools from optimization, statistics, and control theory to better harness flexible loads for battery services.

The recent FERC order brings additional motivation for focus on flexible loads: The introduction of [15] states that FERC order 2222 enables DERs to play a wide range of roles in organized electricity markets, including capacity, energy and ancillary services; moreover, "*multiple DERs can aggregate to satisfy minimum size and performance requirements that they might not meet individually.*" The value of aggregation is clear to FERC and the research community.

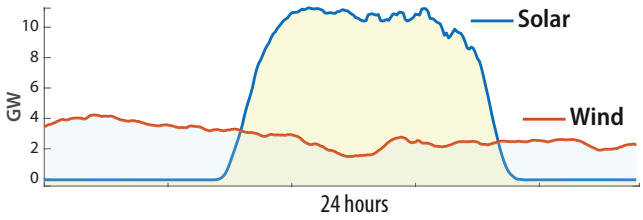

Fig. 1. Solar and Wind in CAISO (California Independent System Operator) on March 16, 2021. Figure generated using data from [1].

A. Demand Dispatch

The term *demand dispatch* refers to control of flexible loads for the creation of *virtual energy storage* (VES) — see [4], [7], [11], [17], [34] and the references therein. Automation is central to the success of demand dispatch: through distributed control, deferrable loads, such as electric water heaters, refrigerators, ACs, etc., provide grid balancing and regulation service while also ensuring consumer-side quality of service (QoS).

The grid services of interest might include ramping, reserves, voltage support, automatic generation control, and peak-shaving. The definition of QoS depends on load-class. For example, a residential refrigerator must keep its content within temperature limits. Any control solution must respect these bounds, along with cycling constraints [14].

Fig. 2 illustrates the communication and control architecture considered in this paper, featuring the following entities.

(i) *Balancing authority* (BA): the entity responsible for the reliable operation of the grid within a balancing authority area (collection of generation, transmission, and loads within a metered boundary). The BA, (also called the *grid operator*), corresponds to an Independent System Operator (ISO) or a Regional Transmission Organization (RTO) in many parts of the U.S., or a large utility in vertically integrated regions (e.g., FP&L and Duke in Florida). It maintains real-time system information such as grid frequency, tie line errors, etc., alongside forecasts of demand and renewable generation.

(ii) *Resource Aggregators*: a single resource aggregator (RA) will receive a power deviation command signal from the grid operator based on current and forecast grid data, and will broadcast a common signal $\{\lambda_t^\star\}$ to each load in its territory.

An RA periodically measures the state of charge (SoC) and capacity of the load aggregations under their purview, and

*Funding from the National Science Foundation under award EPCN 1935389 is gratefully acknowledged. Simon's Institute, Berkeley, is gratefully acknowledged for hosting SM and JW in Spring 2018: discussions during the semester-long program on *Real Time Decision Making* served as inspiration for this paper.

J. Mathias is with the Department of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85281, USA (email: joel.mathias@asu.edu).

S. Meyn is with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA (email: meyn@ece.ufl.edu).

R. Moye is with Tyr Energy, Inc., Ponte Vedra Beach, FL 32081, USA (email: bmoye@tyrenergy.com).

J. Warrington is with AstraZeneca, Cambridge, UK (email:joe.warrington@gmail.com).

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2023.3293037

may share part of this information with the grid operator. Only aggregate-level information is required from the loads in this and prior work, which can be estimated by periodically sampling data from the population [10], [31]. The information is used to obtain initial conditions for optimization, as well as to update VES models.

Examples of RAs in North America include Enbala Power Networks, Comverge, Enel X, and CPower. Utilities can also serve as resource aggregators: examples include Florida Power and Light's *On Call* program and Austin Energy's *Power Partner* program. The former engages air conditioning, electric water heaters, and pool pumps of nearly 800,000 customers, while the latter controls NEST thermostats of participating customers during grid events such as summer peaks and system emergencies.

(iii) *Load classes*: shown in Fig. 2 are M load classes. The first load class consists of a collection of refrigerators, the second corresponds to an aggregation of water heaters, and so forth.

Flexible loads include both residential and commercial thermostatically controlled loads (TCLs), residential pool pumps, and water pumping for irrigation or waste management. There may be hundreds of thousands of loads within each class.

Associated with a given load class, indexed by $i \in \{1, ..., M\}$, is the SoC of the load class, denoted $x_i \in \mathbb{R}$. In the standard VES model, $x_i(t)$, which is by definition proportional to the average QoS of the ith load aggregation at time t, evolves as the state process of a linear system, whose input is the power deviation $-z_i$ [21]. This paper builds upon this model with the introduction of a cost function $c_i : \mathbb{R} \to \mathbb{R}_+$ that penalizes deviation from the ideal QoS. Details are deferred to Section II-A and Section II-B.

Fig. 2 suggests two resource allocation problems: The first is solved by the grid operator, who engages RAs and other entities providing grid services, and will optimally determine the power trajectory requested from each based on available grid data, net-load[1] forecasts, and SoC for both traditional energy storage and VES.

This paper focuses on the second resource allocation problem, which is of concern to an RA that must manage a highly diverse population of DERs. Homogeneity of load behavior within a class is assumed; this may be firmly justified through the introduction of additional local control [31], [32], along with one of the distributed control approaches advocated in [7], [11], [17], [34], [4].

It is shown that this resource allocation problem, formulated as an optimal control problem in Section II, belongs to a special class known as "cheap optimal control" [42], [23]. As a consequence, we will see that the optimal state evolves on a low-dimensional manifold [44], [20].

B. Related research

The past decade has seen rapid development of academic research on the real-time control of flexible loads to provide grid services; several valuable surveys may be found in [37].

[1]net-load is power demand − power generation from renewables.

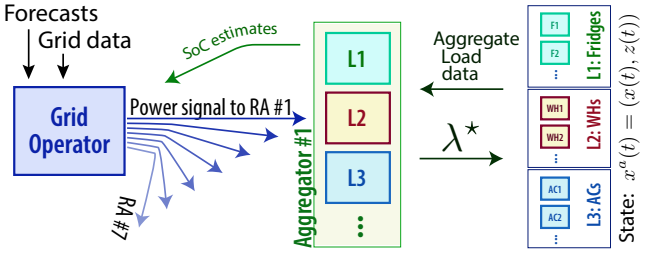

Fig. 2. A distributed control architecture for RA scheduling. A topic addressed in this paper: the interpretation and computation of $\lambda^\star$.

The modeling of large load aggregations was first tackled in the 1980s: in [29], a non-linear model for a large collection of TCLs is obtained using concepts from statistical physics.

The present work was motivated in large part by [4], in which a frequency decomposition was proposed for resource allocation. There is evidence that this approach will indirectly reduce impact on loads. For example, refrigerators have a cycle time of about one hour, so are ideally suited to regulation near this frequency band; pool cleaning is ideal for much lower frequencies. Unfortunately, this approach appears to be conservative. In particular, any load can be turned off, which is inherently "high frequency".

Recent research on the optimal allocation of load aggregations has focused on numerical methods. The paper [5] applies Markov decision process techniques to address the AC optimal power flow problem with participation of TCLs. In [18], the discrete-time resource allocation problem is formulated as a sequence of optimization problems corresponding to each time instance. These optimization problems rely on measurements at each time instance instead of dynamical models of the load classes; their solution defines a command signal for each load class. The optimization formulation presented here is far more general, as is the goal: to reveal the structure of the optimal solution, and use it to inform control design.

The term *state space collapse* in the title comes from the literature on stochastic networks [40], [38], which may be regarded as a special case of the model reduction obtained using singular perturbation methods [42].

In the infinite horizon setting with quadratic cost, the main result of [20] implies that the state space collapse is just one-dimensional: there is a one-dimensional subspace $\mathsf{X}^\star$ such that $x^\star(t) \in \mathsf{X}^\star$ for all $t > 0$. These results are extended to include the finite-time horizon optimal control problem in [22] (see in particular [22, Theorem 5.8]).

Portions of the theory surveyed here appeared in the conference paper [33] and dissertation [30].

C. Contributions

The contributions are adumbrated below.

1) State space collapse: The main technical conclusion is that the optimal state trajectory $x^\star$ evolves on a two-dimensional manifold. This is most conveniently described through a change of variables: the *marginal cost* evolves on

a two-dimensional subspace generated by a scalar co-state process $\lambda^\star$ and its derivative:

$$c_i'\left(x_i^\star(t)\right) = \alpha_i \lambda^\star(t) - \tfrac{d}{dt}\lambda^\star(t)\,, \quad t>0\,, \tag{1}$$

in which $c_i' = dc_i/dx$. This identity justifies the simple control architecture shown in Fig. 2, in which a common signal $\lambda^\star$ is broadcast to all loads, independent of class.

To the best of our knowledge, state space collapse has never been identified in the economic dispatch literature.

Three approaches are used in analysis leading to (1) and its extensions: Bellman's principle of optimality quickly reveals the emergence of state space collapse. Pontryagin's minimum principle is the most direct approach to construct $\lambda^\star$, where it plays the role of a co-state variable. Finally, a Lagrangian relaxation leads to an alternative interpretation of the co-state variable, which will likely provide inspiration for efficient numerical techniques.

The representations of the value functions in terms of lower dimensional variables (see Prop. II.1), and the proofs utilized to arrive at this result (see Appendix A) are a novel contribution to the theory of singular optimal control.

2) Structure for quadratic cost: Section IV-A is devoted to this special case. It is shown that the solution to the $2M$-dimensional control problem may be obtained via the solution of a 2-dimensional differential Riccati equation.

3) Numerical examples: The theory is illustrated with a survey of results from numerical experiments, focusing primarily on managing the grid in California.

D. Organization

Preliminaries required for analysis are introduced in Section II; in particular, the dynamic control problem is introduced in Section II-C. The main results surveyed in Section III demonstrate state space collapse and its consequences. Section IV provides examples: special structure in the case of quadratic cost, and results form simulation studies to illustrate application in realistic settings. Conclusions and directions for future research are contained in Section V, and technical results are contained in the Appendix

Notation

$\mathcal{T} \ge 1$: time horizon for control.
$\ell(t)$: net-load on $[0,\mathcal{T}]$.
$g(t)$: power from traditional generation; $\gamma(t) = \frac{d}{dt}g(t)$.
M: number of load classes, indexed by $i \in \{1,...,M\}$.
$x_i(t)$: state of charge (SoC) of load class i.
$-z_i(t)$: power deviation from load class i; $u_i(t) = \frac{d}{dt}z_i(t)$.
Subscript "σ" denotes sum, e.g., $x_\sigma(t) = \sum_i x_i(t)$.
$x^a := (x,z)$ and $x^a_\sigma := (x_\sigma, z_\sigma)$ are the augmented state and descriptor state, respectively. Consequently, evaluation of a function $f : \mathbb{R}^M \times \mathbb{R}^M \to \mathbb{R}$ is equivalently expressed $f(x,z)$ or $f(x^a)$.
A bar above a variable denotes average either over t or i, e.g.,

$$\bar{\ell} = \frac{1}{\mathcal{T}}\int_0^{\mathcal{T}} \ell(t)\,dt, \qquad \bar{\alpha} = \frac{1}{M}\sum_i \alpha_i\,.$$

Standard calculus notation is adopted for a function $K\colon \mathbb{R}^3 \to \mathbb{R}$:

$$K_\xi(r,s,t) := \left(\tfrac{\partial}{\partial r}\right)^{\xi_1}\left(\tfrac{\partial}{\partial s}\right)^{\xi_2}\left(\tfrac{\partial}{\partial t}\right)^{\xi_3} K(r,s,t)\,, \qquad \xi \in \mathbb{Z}_+^3.$$

$P_t(i,j)$ denotes the (i,j)-th entry of the time-varying matrix P_t. The transpose of a matrix (or vector) P is denoted $P^\intercal$.

I is the identity matrix, E is the matrix of all 1s, $\mathbf{1}$ is the column vector of all 1s, and $\mathbf{0}$ is the column vector or matrix of all 0s. Consequently,

$$\begin{bmatrix} x_\sigma \\ z_\sigma \end{bmatrix} = w^\intercal \begin{bmatrix} x \\ z \end{bmatrix}, \qquad w = \begin{bmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \tag{2}$$

The dynamical system (4) is alternately represented

$$\begin{aligned} &\tfrac{d}{dt}\begin{bmatrix} x(t) \\ z(t) \end{bmatrix} = A\begin{bmatrix} x(t) \\ z(t) \end{bmatrix} + Bu(t) \\ &\textit{where,}\quad A = \begin{bmatrix} -\mathrm{diag}(\alpha) & -I \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \qquad B = \begin{bmatrix} \mathbf{0} \\ I \end{bmatrix} \end{aligned} \tag{3}$$

Parentheses are used for time indices, and subscripts are used to enumerate resource aggregations.

II. Architecture for Resource Allocation

The optimal control architecture for scheduling flexible loads and other resources to meet forecast net-load.

We begin by specifying the dynamic models used for load aggregations.

A. VES models

An aggregate model for dynamics is required for each load class. In the case of TCLs, we adopt the model of [21], and for residential pools or irrigation, we consider the similar model of [9], [36]. In either case, the model for the ith load class is the first-order ODE,

$$\tfrac{d}{dt}x_i(t) = -\alpha_i x_i(t) - z_i(t), \qquad i \in \{1,...,M\}\,, \tag{4}$$

in which $x_i(t)$ is the SoC of the ith load class, and α_i is a non-negative leakage parameter (e.g., for TCLs, this corresponds to the thermal time constant). The power deviation at time t is $-z_i(t)$: this is the total power consumed by the load aggregation minus the nominal power consumption of the loads (nominal refers to the uncontrolled behavior of the load aggregation, i.e., the power consumption when the loads are not participating in demand dispatch). Consequently, $z_i(t)$ is the (virtual) power supplied by the VES to grid at time t from load class i.

The precise adjustments in power consumption are obtained as the solution to an optimal control problem. A lower level control problem is required to ensure that the population of loads in class i collectively adjust their power consumption to approximately track the optimal trajectory $z_i^\star$. This paper is agnostic to the precise way in which control is employed to achieve tracking. Many approaches are described in prior work—see recent surveys in [37].

B. Cost formulation and state space collapse

The resource allocation problem is formulated as a finite-horizon optimal control problem over $[0, \mathcal{T}]$, with objective

$$\int_0^{\mathcal{T}} \left[c_{\mathrm{g}}(g(t)) + c_{\mathrm{d}}(\tfrac{d}{dt} g(t)) + c_{\mathrm{x}}(x(t)) \right] dt \tag{5}$$

in which the cost functions c_{x}, c_{g} and c_{d} are strongly convex. The state cost is expressed as a sum:

$$c_{\mathrm{x}}(x) = \sum_{i=1}^{M} c_i(x_i)\,, \quad x \in \mathbb{R}^M\,, \tag{6}$$

where each $c_i \colon \mathbb{R} \to \mathbb{R}_+$ is a strongly convex function, designed to impose high penalty when the SoC strays outside of desired capacity bounds C_i, as illustrated in Fig. 3.

There are good reasons to avoid hard constraints on SoC or generation, even though there are limits in practice. We would lose much of the elegant structure with a state constrained optimization problem, and we do not see much benefit. Remember, the optimization problems posed here will *not* be implemented in real time. It is anticipated that the solution will be part of a model predictive control (MPC) architecture. If severe constraint violations are observed in offline computations, then the cost functions may be modified to obtain a feasible solution, or additional resources may be required to meet constraints.

We could instead opt for a *barrier function* approach in which the cost function shown in Fig. 3 is infinite outside of an open interval. Theory in this paper can be extended to this setting without significant changes. The choice of a penalty function is to improve numerical stability in simulations. We have found in experiments that constraint violations are negligible for costs of the form Fig. 3 and the reference signals scaled to ensure feasibility.

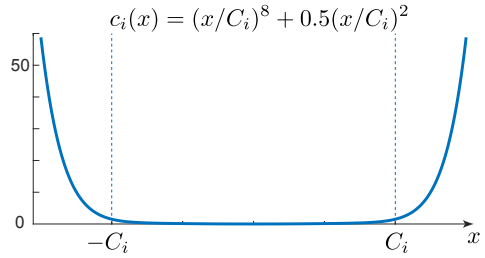

Fig. 3. Cost of QoS violation

Generation costs are translated to costs on power deviation through the supply-demand constraint:

$$\ell(t) = g(t) + z_\sigma(t)\,, \qquad 0 \le t \le \mathcal{T} \tag{7}$$

in which ℓ denotes forecast net-load, and the subscript denotes summation: $z_\sigma(t) = \sum_i z_i(t)$. Because a cost is imposed on both $g(t)$ and its derivative, to put the optimal control problem in standard form requires state augmentation. The augmented state is denoted $x^a = (x, z) \in \mathbb{R}^{2M}$, with M-dimensional input $u(t) = \frac{d}{dt} z(t)$, $t \ge 0$.

The source of collapse. The optimal control formulations considered in this paper fall in the category of *cheap optimal control*. This conclusion is reached by expressing the generation and its derivative at time t as

$$g(t) = \ell(t) - z_\sigma(t) \qquad \tfrac{d}{dt} g(t) = \tfrac{d}{dt} \ell(t) - u_\sigma(t)$$

It follows that the objective (5) imposes a cost on the *sums* $z_\sigma(t)$ and $u_\sigma(t)$, and in particular the cost on the input $u(t) \in \mathbb{R}^M$ will not be coercive. Basic theory predicts that there is a low-dimensional dynamical systems representation of the optimal control solution, based on the *descriptor state* [22], [42].

It is shown in the present paper that the descriptor dynamics take the following form,

$$\tfrac{d}{dt} x_\sigma(t) = -\alpha^\intercal x(t) - z_\sigma(t), \tag{8a}$$

$$\tfrac{d}{dt} z_\sigma(t) = u_\sigma(t). \tag{8b}$$

where again the subscripts indicate sums, and the components of the column vector $\alpha \in \mathbb{R}^M$ are the linear system coefficients, α_i, appearing in (4).

C. Optimal control

The optimal control problem over the finite time-horizon $[0, \mathcal{T}]$ is defined as follows: with $x(0), z(0) \in \mathbb{R}^M$, given,

$$\underset{g,\gamma,x,z,u}{\text{minimize}} \quad \int_0^{\mathcal{T}} \left[c_{\mathrm{g}}(g(t)) + c_{\mathrm{d}}(\gamma(t)) + c_{\mathrm{x}}(x(t)) \right] dt \tag{9a}$$

$$\text{subject to} \quad \ell(t) = g(t) + z_\sigma(t) \quad , \tag{9b}$$

$$\tfrac{d}{dt} g(t) = \gamma(t) \quad , \tag{9c}$$

$$\tfrac{d}{dt} x_i(t) = -\alpha_i x_i(t) - z_i(t) \quad , \tag{9d}$$

$$\tfrac{d}{dt} z_i(t) = u_i(t), \quad i \in \{1, ..., M\} \tag{9e}$$

The objective (9a) is the total cost of supplying the net load ℓ using traditional generation and VES from the loads participating in demand dispatch. Consequently, there are costs on: (i) traditional generation, modeled through c_{g}; (ii) generation ramping, modeled through c_{d}; and (iii) SoC of the load classes, modeled through c_{x}. The constraint (9b) is the demand-supply balancing constraint, i.e., the net load ℓ is balanced using the output of the traditional generation g and power deviation from flexible loads z. The equations (9d) and (9e) provide the dynamical constraints on load classes.

The analysis here allows general strongly convex and twice continuously differentiable cost on SoC and generation, but imposes a quadratic cost on ramping.

Assumption 1. *The net load ℓ is C^1 on $[0, \mathcal{T}]$, and its derivative is Lipschitz continuous. The cost functions $\{c_i\}$ and c_g are non-negative, class C^2, and strongly convex: $c_g''(x) \ge \mu$ and $c_i''(x) \ge \mu$ for some $\mu > 0$ and all i, x. The ramping cost is quadratic: for fixed $\kappa > 0$,*

$$c_d(x) = \tfrac{1}{2}\kappa x^2, \qquad x \in \mathbb{R}. \tag{10}$$

For notational convenience, the cost beyond the control time horizon is taken to be zero: $c(x, z, u, t) := 0$ for all $t > \mathcal{T}$, and triple (x, z, u).

The optimization variables g and γ may be eliminated by applying the algebraic constraint (9b). The overall optimization problem (9a) is then in the standard form used in optimal

control textbooks, with state process is $x^a(t) = (x(t), z(t)) \in \mathbb{R}^{2M}$ and cost function

$$c(x^a(t), u(t), t) = c_{\mathrm{x}}(x(t)) + c_{\mathrm{g}}(\ell(t) - z_\sigma(t)) + c_{\tilde{d}}(u(t), t) \tag{11a}$$

$$\textit{where} \quad c_{\tilde{d}}(u(t), t) = \tfrac{1}{2}\kappa\big[u_\sigma(t) - \tfrac{d}{dt}\ell(t)\big]^2 \tag{11b}$$

The total cost in (9a) is the integral of (11a), from which cost degeneracy is again evident: the terms involving the control cost in (11a) are expressed purely in terms of the sum $u_\sigma(t)$. Implications are described in the following, beginning with properties of the value function.

D. Value functions

For $t_0 \in [0, \mathcal{T})$, the cost-to-go is denoted,

$$J^\star(x, z, t_0) := \inf_{u_{[t_0,\mathcal{T}]}} \int_{t_0}^{\mathcal{T}} c(x^a(t), u(t), t)\, dt, \tag{12}$$

where the infimum is over continuous u, subject to (9d)–(9e), and with $x(t_0) = x, z(t_0) = z$ given.

Prop. II.1 asserts that the cost-to-go can be expressed purely as a function of $x^a_\sigma = (x_\sigma, z_\sigma)$. This is the first evidence of state space collapse. Denote

$$\begin{aligned} &K^\star(x_\sigma, z_\sigma, t_0) := \inf_{x^+, z^+} J^\star(x^+, z^+, t_0), \\ &\text{subject to} \qquad x^+_\sigma = x_\sigma, \qquad z^+_\sigma = z_\sigma, \end{aligned} \tag{13}$$

where the infimum is over $x^+, z^+ \in \mathbb{R}^M$.

The proof of Prop. II.1 and most of the results that follow are contained in the Appendix.

Proposition II.1. *The following hold under Assumption 1: for each $t_0 \in [0, \mathcal{T})$,*

(i) *$J^\star$ is convex in $x^a = (x, z)$ and finite-valued.*

(ii) *$J^\star(x, z, t_0) = K^\star(x_\sigma, z_\sigma, t_0)$ for each $x, z \in \mathbb{R}^M$.*

For a given initial condition x, z, and resulting optimal state trajectory $\{x^\star(t), z^\star(t) : 0 < t \le \mathcal{T}\}$, denote

$$\begin{aligned} \lambda^\star(t) &= K^\star_{1,0,0}(x^\star_\sigma(t), z^\star_\sigma(t), t), \\ \beta^\star(t) &= K^\star_{0,1,0}(x^\star_\sigma(t), z^\star_\sigma(t), t)\,, \qquad 0 < t \le \mathcal{T}. \end{aligned} \tag{14}$$

In addition to Assumption 1, the following assumptions are imposed throughout the remainder of the paper:

Assumption 2. *For each $t_0 \in [0, \mathcal{T})$ and each initial condition (x, z), the optimal control problem admits a unique solution $\{x^\star(t), z^\star(t), u^\star(t) : t_0 < t \le \mathcal{T}\}$ satisfying the following:*

(a) $(x^\star(t), z^\star(t))$ is C^1 on the semi-open interval $(0, \mathcal{T}]$.

(b) There are right hand limits at t_0, denoted

$$x^\star(t_0+) = \lim_{t \downarrow t_0} x^\star(t)\,, \qquad z^\star(t_0+) = \lim_{t \downarrow t_0} z^\star(t), \tag{15}$$

satisfying $x^\star_\sigma(t_0+) = x_\sigma$, $z^\star_\sigma(t_0+) = z_\sigma$.

Assumption 3. *The value function $K^\star\colon \mathbb{R} \times \mathbb{R} \times [0, \mathcal{T}] \to \mathbb{R}$ is C^1.*

Assumption 4. *The functions $\lambda^\star, \beta^\star$ are C^1.*

It is shown in Section IV-A that Assumptions 2 to 4 hold for quadratic cost. We have not found a reference to justify these assumptions for more general cost functions, but the algorithms we have tested yield results that are consistent with the general theory surveyed in this paper. Some examples are described in Section IV-B.

III. State space collapse

Thm. III.1 unveils the structure of the optimal solution. In particular, the M-dimensional optimal state process $x^\star$ evolves on a two-dimensional manifold.

Theorem III.1. *The optimal solution $(x^\star, z^\star, u^\star, \lambda^\star, \beta^\star)$ is the solution to the following system of $2M + 2$ equations:*

$$\tfrac{d}{dt}x^\star_i(t) = -\alpha_i x^\star_i(t) - z^\star_i(t), \tag{16a}$$

$$\tfrac{d}{dt}z^\star_i(t) = u^\star_i(t), \tag{16b}$$

$$\tfrac{d}{dt}\lambda^\star(t) = -c_i'\left(x^\star_i(t)\right) + \alpha_i \lambda^\star(t), \qquad \textit{for each} \quad i \in \{1, ..., M\}, \tag{16c}$$

$$\tfrac{d}{dt}\beta^\star(t) = c_g'(\ell(t) - z^\star_\sigma(t)) + \lambda^\star(t), \tag{16d}$$

$$u^\star_\sigma(t) = \tfrac{d}{dt}\ell(t) - \frac{1}{\kappa}\beta^\star(t), \tag{16e}$$

with boundary conditions $\lambda^\star(\mathcal{T}) = \beta^\star(\mathcal{T}) = 0$, and $x^\star(0+), z^\star(0+)$ defined in Assumption 2.

Three different interpretations of $\lambda^\star$ are explained in the following.

A. λ^* as the command signal

Equation (16c) has a remarkable interpretation: the marginal costs for the M different load classes evolve in a two-dimensional subspace generated by the functions $\{\lambda^\star(t), \frac{d}{dt}\lambda^\star(t) : t \in (0, \mathcal{T}]\}$.

Since c_i is strictly convex, c_i' is strictly monotone and bijective. Consequently, the optimal SoC evolves on a two-dimensional manifold and can be computed based on $\lambda^\star$ and its derivative:

$$x^\star_i(t) = (c_i')^{-1}(\alpha_i \lambda^\star(t) - \tfrac{d}{dt}\lambda^\star(t)). \tag{17}$$

B. $\lambda^\star$ as co-state

The Hamiltonian with co-state variables $\lambda, \beta \in \mathbb{R}^M$ corresponding to system equations (9d) and (9e), respectively, is denoted:

$$H(x, z, u, \lambda, \beta, t) := c(x, z, u, t) + \sum_i \lambda_i(-\alpha_i x_i - z_i) + \sum_i \beta_i u_i \tag{18}$$

This notation for the co-state variables may appear to conflict with the notation in Thm. III.1. The choice of notation is made clear in the following:

Proposition III.2. *Associated with the optimal input-state $(u^\star, x^\star, z^\star)$ are a pair of co-state variables $\lambda^\star, \beta^\star$ evolving in $\mathbb{R}^M$ and satisfying for $0 < t \le \mathcal{T}$,*

$$(x^\star(t), z^\star(t)) = \arg\min_{(x,z)} H(x, z, u^\star(t), \lambda^\star(t), \beta^\star(t), t)$$

For each $i \in \{1, ..., M\}$ and $t \in (0, \mathcal{T}]$,

$$\begin{aligned} \lambda_i^\star(t) &= K^\star_{1,0,0}(x^\star_\sigma(t), z^\star_\sigma(t), t) = \lambda^\star(t), \\ \beta_i^\star(t) &= K^\star_{0,1,0}(x^\star_\sigma(t), z^\star_\sigma(t), t) = \beta^\star(t). \end{aligned} \tag{19}$$

Proof. The left hand equalities in (19) are a familiar result: the optimal co-state trajectory is the gradient of the value function with respect to the state variable [13, Theorem 3.1]. The right-hand equalities are from the definition (14). ■

C. $\lambda^\star$ and a Lagrangian decomposition

Rather than eliminate the variable g using (9b), new insight is obtained on maintaining g, x as variables in the optimization problem. First, construct a Lagrangian relaxation with Lagrange multiplier ϱ corresponding to the algebraic constraint (9b), as follows:

$$\begin{aligned} \phi^\star(\varrho) = \inf_{g,x} \int_0^{\mathcal{T}} \Big\{ & c_{\mathrm{g}}(g(t)) + c_{\mathrm{d}}(\dot g(t)) + c_{\mathrm{x}}(x(t)) \\ & + \varrho(t)(\ell(t) - g(t) - z_\sigma(t)) \Big\}\, dt, \end{aligned} \tag{20}$$

where the infimum is subject to (9d)–(9e), with given initial conditions.

This amounts to a Lagrangian decomposition, consisting of the following $M + 1$ independent optimization problems:

(i) Minimization problem over $\{g(t), \dot g(t)\}$:

$$\inf_g \int_0^{\mathcal{T}} \mathcal{L}_g(g(t), \dot g(t), t)\, dt, \tag{21}$$

where,

$$\mathcal{L}_g(g(t), \dot g(t), t) = c_{\mathrm{g}}(g(t)) + c_{\mathrm{d}}(\dot g(t)) - \varrho(t)(g(t) - \ell(t)).$$

(ii) Minimization problem over $\{x_i(t), \dot x_i(t)\}$:

$$\inf_{x_i} \int_0^{\mathcal{T}} \mathcal{L}_i(x_i(t), \dot x_i(t), t)\, dt, \tag{22}$$

where, after accounting for the constraints (9d), (9e),

$$\mathcal{L}_i(x_i(t), \dot x_i(t), t) = c_i\big(x_i(t)\big) + \alpha_i \varrho(t) x_i(t) + \varrho(t) \dot x_i(t).$$

The Euler-Lagrange equations lead to equations for the optimizers:

Proposition III.3. *For any function ϱ that is continuously differentiable on $(0, \mathcal{T}]$, if g^ϱ and x_i^ϱ are C^1 optimizers for the minimization problems in eqs.* (21) *and* (22)*, then they solve the following differential equations:*

$$c_g'\left(g^\varrho(t)\right) - \tfrac{d}{dt} c_d'(\dot g^\varrho(t)) = \varrho(t), \tag{23}$$

$$c_i'\left(x_i^\varrho(t)\right) + \alpha_i \varrho(t) - \tfrac{d}{dt}\varrho(t) = 0, \tag{24}$$

with boundary conditions $c_d'(\dot g^\varrho(\mathcal{T})) = \varrho(\mathcal{T}) = 0$.

Proof. As $\varrho \in C^1$, it follows from Assumption 1 that $\mathcal{L}_g, \mathcal{L}_i \in C^1$. Moreover, $(g^\varrho(t), \dot g^\varrho(t))$ and $(x_i^\varrho(t), \dot x_i^\varrho(t))$ are continuous on the half-open interval $(0, \mathcal{T}]$. Consequently, the Euler-Lagrange equations form the necessary first-order conditions for weak extrema [26, Section 2.3.3]. The solution to the minimization problem in (21) and (22) at the stationary minimum are the following Euler-Lagrange equations,

$$\begin{aligned} \tfrac{\partial}{\partial g}\mathcal{L}_g(g^\varrho, \dot g^\varrho, t) - \tfrac{d}{dt}\tfrac{\partial}{\partial \dot g}\mathcal{L}_g(g^\varrho, \dot g^\varrho, t) = 0, \\ \tfrac{\partial}{\partial x_i}\mathcal{L}_i(x_i^\varrho, \dot x_i^\varrho, t) - \tfrac{d}{dt}\tfrac{\partial}{\partial \dot x_i}\mathcal{L}_i(x_i^\varrho, \dot x_i^\varrho, t) = 0, \end{aligned}$$

which result in (23) and (24), respectively. The terminal-time boundary conditions are obtained by, respectively, setting $\tfrac{\partial}{\partial \dot g}\mathcal{L}_g(g^\varrho, \dot g^\varrho, t)|_{t=\mathcal{T}} = 0$ and $\tfrac{\partial}{\partial \dot x_i}\mathcal{L}_i(x_i^\varrho, \dot x_i^\varrho, t)|_{t=\mathcal{T}} = 0$ [26, Section 2.3.5]. ■

The dual functional $\phi^\star$ satisfies weak duality: $\phi^\star(\varrho) \le J^\star(x(0), z(0), 0)$ for any ϱ, and the dual convex program is defined as $\sup_\rho \phi^\star(\varrho)$. The solution to the dual is obtained by combining Prop. III.3 and Thm. III.1, and from this we obtain strong duality:

Proposition III.4. *The dual admits an optimizer given by*

$$\varrho^\star(t) = -\lambda^\star(t)\,, \qquad t \in (0, \mathcal{T}].$$

Proof. With $t \in (0, \mathcal{T}]$, setting $\varrho(t) = -\lambda^\star(t)$ in (24) and comparing with (16c) yields $x_i^\varrho(t) = x_i^\star(t)$ for each i, which is the optimal solution to the primal problem (9a). This implies that there is no duality gap: $-\lambda^\star$ maximizes $\phi^\star$. ■

IV. Examples

We conclude with two general examples: Section IV-A describes the optimal solution for the special case of quadratic cost, and Section IV-B provides results from numerical experiments for the general, polynomial cost functions like the ones shown in Fig. 3.

A. Quadratic cost functions and the Riccati equation

It is instructive to compute the optimal solutions for the special case of quadratic cost functions, and constant $\ell(t) \equiv \overline{\ell}$.

The quadratic costs are given as follows:

$$\begin{aligned} c_i(x) &= \tfrac{1}{2} q_i x^2, && q_i > 0, \\ c_{\mathrm{g}}(\ell(t) - z_\sigma(t)) &= \tfrac{1}{2} q_z (z_\sigma(t))^2, && q_z > 0 \\ c_{\tilde d}(u(t)) &= \tfrac{1}{2}\kappa (u_\sigma(t))^2 \end{aligned} \tag{25}$$

The second equation in (25) implies a quadratic cost on generation, given by $c_{\mathrm{g}}(g(t)) = \tfrac{1}{2} q_z (g(t) - \overline{\ell})^2$.

The value function of interest is again,

$$J^\star(x^a, t_0) = \inf_{u[t_0, \mathcal{T}]} \int_{t_0}^{\mathcal{T}} c(x^a(t), u(t), t)\, dt \tag{26}$$

subject to $x^a(t_0) = x^a$ and the constraints (9d), (9e), with

$$c(x^a(t), u(t), t) = \tfrac{1}{2} x^a(t)^{\intercal} Q x^a(t) + c_{\tilde d}(u(t), t) \tag{27}$$

where

$$Q = \begin{bmatrix} \mathrm{diag}(q_i) & \mathbf{0} \\ \mathbf{0} & q_z E \end{bmatrix}$$

and E is the matrix with all entries equal to unity.

The value function $K^\star$ admits a representation in terms of the solution to a 2×2 matrix differential Riccati equation.

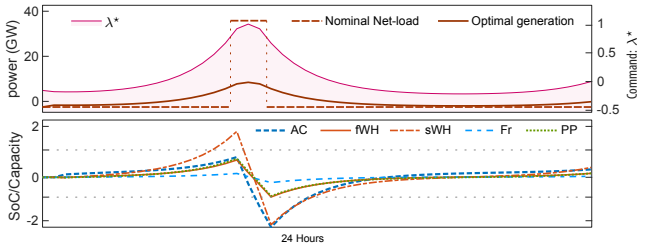

Fig. 4. Optimal solution with all costs quadratic. The optimal generation $g^\star$ is nearly constant, despite the 40 GW load surge.

Proposition IV.1. *The value function* (26) *may be expressed,*

$$J^\star(x^a,t) = K^\star(x^a_\sigma,t) = \tfrac{1}{2}(x^a_\sigma)^\intercal P_t x^a_\sigma \tag{28}$$

where $P_t \geq 0$ is the solution to the differential Riccati equation,

$$\begin{aligned} 0 = \dot{P}_t + P_t \begin{bmatrix} -\bar{\alpha} & -1 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} -\bar{\alpha} & 0 \\ -1 & 0 \end{bmatrix} P_t \\ + \begin{bmatrix} \frac{\bar{q}}{M} & 0 \\ 0 & q_z \end{bmatrix} - \frac{1}{\kappa} P_t \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} P_t \end{aligned} \tag{29}$$

with boundary condition $P_{\mathcal{T}} = 0$.

The evolution of the "projected" optimal state trajectory $x^{a\star}_\sigma(t)$ can be represented using P_t, which consequently defines a map between the initial state and the optimal subspace, as shown in the following proposition.

Proposition IV.2. (i) *The evolution of $x^{a\star}_\sigma(t)$ is given by the linear time-varying system:*

$$\tfrac{d}{dt} x^{a\star}_\sigma(t) = A^\sigma_t x^{a\star}_\sigma(t), \qquad t \in [0,\mathcal{T}], \tag{30}$$

with initial condition $x^{a\star}_\sigma(0) = [x_\sigma(0), z_\sigma(0)]^\intercal$, and $A^\sigma_t = S_t^{-1} T_t$ with,

$$S_t = \begin{bmatrix} \xi_1 P_t(1,1) & \xi_1 P_t(1,2) \\ 0 & 1 \end{bmatrix}$$

$$T_t = \begin{bmatrix} \xi_2(P_t(1,1) - \dot{P}_t(1,1)) & \xi_2(P_t(1,2) - \dot{P}_t(1,2) - 1) \\ -\kappa^{-1} P_t(2,1) & -\kappa^{-1} P_t(2,2) \end{bmatrix}$$

$$\xi_1 = 1 - \sum_i \frac{\alpha_i}{q_i} \qquad \xi_2 = -\sum_i (1-\alpha_i)\frac{\alpha_i}{q_i}$$

(ii) *There exist maps $G^1 : \mathbb{R}^{2\times 2} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^M$ and $G^2 : \mathbb{R}^{2\times 2} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^M$ such that*

$$\begin{aligned} x^\star(t) &= G^1(P_t, x^a_\sigma(0)) \\ z^\star(t) &= G^2(P_t, x^a_\sigma(0)), \quad t \in (0,\mathcal{T}] \end{aligned} \tag{31}$$

The right hand side of each equation is linear in $x^a_\sigma(0)$.

(iii) *These maps describe the jump location at time $t = 0+$:*

$$x^\star(0+) = G^1(P_0, x^a_\sigma(0)), \quad z^\star(0+) = G^2(P_0, x^a_\sigma(0))$$

Descriptions of the mappings G^1, G^2 may be found in the Appendix, below eq. (47).

B. Demand dispatch - numerical examples

Simulations were conducted to validate the main results of this paper. A discrete-time approximation of the resource allocation problem (9a) was solved with 5 classes of loads: ACs, residential WHs with faster time cycles (fwh), commercial WHs with slower time cycles (swh), refrigerators, and pool pumps (pp), based on the model used in [8].

Load shedding. The first example illustrates the evolution of SoC and the Lagrange multiplier in application to *load shedding*, which is typically the domain of demand response rather than demand dispatch [7].

The net-load is piecewise constant, with a 40 GW surge for 90 minutes in the late morning hours, and the cost on SoC was chosen quadratic for each load class. Results from optimization are collected together in Fig. 4.

The behavior of the loads is similar to that of batteries: they "charge" prior to the outage (i.e., they consume more energy than nominal), which makes it possible to "discharge" (i.e., consume less energy than nominal) during the outage.

The SoC plots reveal that the residential air conditioners and commercial water heaters exceed capacity bounds by a few degrees for a few hours: a SoC violation occurs when |SoC|/Capacity > 1. This is not surprising, given that the solution results in a massive change in net-load.

This motivates the use of more aggressive cost in optimization. In a second set of experiments, the cost functions were chosen to be higher-order polynomials of the form $c_i(x) = \kappa_1(x/C_i)^8 + \kappa_2(x/C_i)^2$, where C_i is the energy capacity of the load class i in GWh. The quadratic term is maintained to ensure strong convexity. It is found that the SoC violations are minor when compared to Fig. 4.

In practice, addressing a 40 GW surge would require more ancillary services than used in these experiments.

Taming the duck. The experiments surveyed next use more realistic day-ahead conditions: the net-load ℓ is based on California's net load ("duck curve") in March 2020, obtained from CAISO.

A common 8th order polynomial cost function was adopted for each class of TCLs, of the form described above with $\kappa_1 = 1$ and $\kappa_2 = 0.1$. Because QoS violations are less important for pool pumps, the quadratic cost was maintained for this load class, using $\kappa_1 = 0$ and $\kappa_2 = 1$. Generation cost was also taken quadratic, of the form $c_{\mathrm{g}}(x) = \kappa_g[x - \overline{\ell}]^2$, where κ_g is a constant gain. Table I provides values of the SoC leakage parameters α_i for the different load classes along with the energy capacities C_i. The numbers are based on [8], [35].

TABLE I: LOAD PARAMETERS

Par.	Unit	ACs	fWHs	sWHs	RFGs	PPs
α_i	hours^{-1}	0.25	0.04	0.01	0.10	0.004
C_i	GWh	4	2	5	0.5	2

The top half of Fig. 5 shows the net-load ℓ, the optimal traditional generation $g^\star$, and the command $\varrho^\star$ (normalized to ± 1 \$/unit energy), while the latter half shows the optimal SoC trajectories normalized by the respective energy capacities, $x^\star_i/C_i$. We observe that $x^\star_i/C_i$ remains within the desired limits ± 1 over the time horizon.

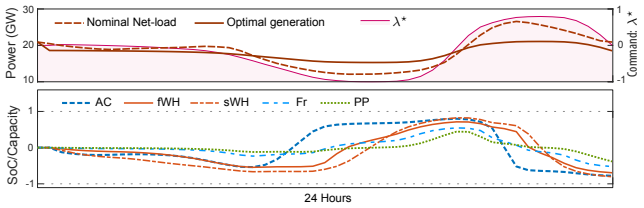

Fig. 5. Optimal SoC trajectories remain within capacity bounds throughout this run. The peak generation for the optimal solution is about 8 GW less than what would be required without load control, and the generation ramping is reduced significantly.

The optimal SoC trajectory for any load class can be recovered based on observations of the SoC for two other load classes by making use of state space collapse. For sake of illustration, suppose we are given the optimal SoC of residential water heaters and ACs. The functions $\{\lambda^\star(t), \frac{d}{dt}\lambda^\star(t) : 0 < t \le \mathcal{T}\}$ are obtained via

$$\begin{bmatrix} \lambda^\star(t) \\ \frac{d}{dt}\lambda^\star(t) \end{bmatrix} = \begin{bmatrix} \alpha_{ac} & -1 \\ \alpha_{fwh} & -1 \end{bmatrix}^{-1} \begin{bmatrix} c_i'(x^\star_{ac}(t)) \\ c_i'(x^\star_{fwh}(t)) \end{bmatrix}$$

We can hence recover any load trajectories using (17).

Fig. 6 shows the SoC trajectory of pool pumps recovered using the trajectories of ACs and residential water heaters.

V. Conclusions and Future Work

The optimal resource allocation problem posed in this paper can be reduced to just two dimensions, regardless of the number of load classes under consideration. The solution can be represented as a distributed control architecture in which a *common* scalar command signal $\lambda^\star$ is broadcast to all the heterogeneous load classes managed by the RA.

State space collapse has valuable computational implications. In particular, Prop. II.1 provides a reduced-dimensional representation of the value function which may be used to inform basis selection in reinforcement learning algorithms. In TD-learning, the basis can be chosen as a function of $x^a_\sigma \in \mathbb{R}^2$ instead of $x^a \in \mathbb{R}^{2M}$, thereby leading to a significant reduction in complexity. Results from preliminary testing of this approach in conjunction with Q-learning may be found in the pre-print [28].

There is, of course, the question of sensitivity to load dynamics and capacity, which will vary with time. It is conjectured that a more robust solution will be obtained in a stochastic control formulation in which disturbances are included in the load dynamics as well as the load forecast. There is a substantial literature on singular control for stochastic systems (e.g., [24]), so this extension may not a great challenge.

Model uncertainty/sensitivity is traditionally addressed using MPC: in the context of this paper, MPC amounts to re-solving (9a) periodically to create an additional layer of feedback.

Optimization and markets. It will be clear to many readers that the resource allocation problem considered in this paper is similar to the starting point of the dynamic competitive equilibrium analysis posed in [12], [39].[2] In this context, this optimization problem may be regarded as the *social planner's problem* (SPP) of micro-economics, and the arguments in [39] together with Prop. III.4 imply that $\varrho^\star$ is the equilibrium price in the dynamic competitive equilibrium model. In this context, Fig. 4 serves as a warning: the "correct" price for a load shedding event is smooth, rather than the surge price typically advocated in the literature on demand response [19].

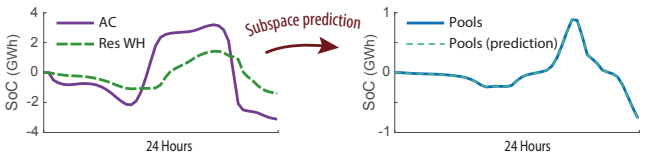

Fig. 6. SoC for pool pumps recovered using those of ACs and WHs.

However, this interpretation ignores many realities: 1) This is a control problem, based on *forecast net-load*. These forecasts will change, perhaps hourly, inducing unpredictable price shocks. 2) Market theory requires some mechanism for price discovery. We cannot see how price discovery is consistent with the control architecture proposed here. 3) Finally, there is the subject of risk and reliability to all involved when the grid is controlled through real time price signals [27], [3], [2].

The markets aspect of demand dispatch remains an open area for research.

Appendix

A. Cheap control and value functions

Lemma A.1. *Consider the dynamical system with state $(x, z) \in \mathbb{R}^{2M}$ and input $u \in \mathbb{R}^M$, with dynamics of specified by* (9d, 9e). *For given $t_0 \ge 0$, suppose that $(x(t_0), z(t_0))$ and (x^+, z^+) are two state values satisfying $x_\sigma(t_0) = x^+_\sigma$, $z_\sigma(t_0) = z^+_\sigma$.*

Then, for each $\delta \in (0, 1)$, there is a C^∞ input u satisfying $u_\sigma(t) = 0$ for all $t_0 \le t \le t_0 + \delta$, and the resulting state trajectory from $x^a(t_0)$ satisfies

$$z(t_0 + \delta) = z^+, \qquad x(t_0 + \delta) = x^+ + O(\delta).$$

Moreover, the following bound holds for $t_0 \le t \le t_0 + \delta$, and $i \in \{1, ..., M\}$:

$$|x_i(t)| \le |x_i(t_0)| + |x_i^+ - x_i(t_0)| + \delta \max(|z_i(t_0)|, |z_i^+|). \quad (32)$$

Proof. Without loss of generality we take $t_0 = 0$. Let $f \colon \mathbb{R} \to \mathbb{R}_+$ be a C^∞ probability density, with support on the interval $(0, \delta)$, and choose

$$u_i(t) = [z_i^+ - z_i(0)]f(t) - [x_i^+ - x_i(0)]f'(t)$$

where f' denotes the derivative of f. This is a "cheap control", since $u_\sigma(t) = 0$ for all t. We then have by definition

$$z_i(t) = z_i(0) + [z_i^+ - z_i(0)] \int_0^t f(\tau)\, d\tau - [x_i^+ - x_i(0)]f(t)$$

[2]These papers, mirroring prior static analyses, began with the assumption that *electric power* is the commodity of interest to the consumers, ignoring the ultimate value of energy to the consumer as in the present work.

This gives $z_i(\delta) = z_i^+$, and

$$\begin{aligned} x_i(t) &= e^{-\alpha_i t} x_i(0) - \int_0^t e^{-\alpha_i (t-\tau)} z_i(\tau)\, d\tau \\ &= x_i(0) - \int_0^t z_i(\tau)\, d\tau + O(\delta) \\ &= x_i(0) + [x_i^+ - x_i(0)] \int_0^t f(\tau)\, d\tau + O(\delta) \end{aligned}$$

which gives $x_i(\delta) = x_i^+ + O(\delta)$.

We have $|z_i(\tau)| \le \max(|z_i(0)|, |z_i^+|) + |x_i^+ - x_i(0)| f(\tau)$ for $0 \le \tau \le \delta$, from which the bound on $x_i(t)$ follows. ■

The following standard result is need in the results that follows—see [41, Theorem A] for a proof.

Lemma A.2. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a convex function. Then, f is locally Lipschitz: that is, for each $x \in \mathbb{R}^n$, there exists a neighborhood U containing x such that f is Lipschitz continuous on U. Equivalently, f is Lipschitz on any compact subset K of $\mathbb{R}^n$.*

The next lemma provides the proof of the first part of Prop. II.1, and further structure on the value function.

Lemma A.3. *The following hold for the value function $J^\star$ under Assumption 1: it is convex in $x^a = (x, z)$, finite-valued, and locally Lipschitz in (x^a, t).*

Proof. It is obvious that $J^\star$ is finite valued. To see that it is convex, let $\{(x^i, z^i) : i = 0, 1\}$ denote two initial conditions (starting at time t_0), fix $\theta \in (0, 1)$, and denote $(x^\theta, z^\theta) = (1-\theta)(x^0, z^0) + \theta(x^1, z^1)$. It remains to show that $J^\star(x^\theta, z^\theta, t_0) \le (1-\theta) J^\star(x^0, z^0, t_0) + \theta J^\star(x^1, z^1, t_0)$ for each t_0. Consider any continuous input-state trajectories:

$$\{(u^i_{[t_0,\mathcal{T}]}, x^i_{[t_0,\mathcal{T}]}, z^i_{[t_0,\mathcal{T}]}) : i = 0, 1\}$$

with given initial conditions $x^i(t_0) = x^i$, $z^i(t_0) = z^i$. Because the system is linear, it follows that the convex combination is feasible from (x^θ, z^θ): with $u^\theta_{[t_0,\mathcal{T}]}$ defined as the convex combination of the inputs, the resulting state trajectory is the convex combination $(x^\theta_{[t_0,\mathcal{T}]}, z^\theta_{[t_0,\mathcal{T}]})$. Consequently,

$$\begin{aligned} J^\star(x^\theta, z^\theta, t_0) &\le \int_{t_0}^{\mathcal{T}} c(x^\theta(t), z^\theta(t), u^\theta(t), t)\, dt \\ &\le (1-\theta) \int_{t_0}^{\mathcal{T}} c(x^0(t), z^0(t), u^0(t), t)\, dt \\ &\quad + \theta \int_{t_0}^{\mathcal{T}} c(x^1(t), z^1(t), u^1(t), t)\, dt, \end{aligned}$$

where the first inequality is the definition of $J^\star$ as an infimum, and the second follows from convexity of the cost function. The proof of convexity is completed on taking the infimum over $u^i_{[t_0,\mathcal{T}]}$ for each $i = 0, 1$.

Since $J^\star$ is convex in x^a, it follows from Lemma A.2 that $J^\star$ is locally Lipschitz in x^a. That is, for each $x^a \in \mathbb{R}^{2M}$, there exists a neighborhood U containing x^a and an $L \ge 0$ such that, for any $t \in [0, \mathcal{T}]$:

$$|J^\star(x_1^a, t) - J^\star(x_2^a, t)| \le L \|x_1^a - x_2^a\|, \qquad x_1^a, x_2^a \in U.$$

Next, we show that $J^\star$ is locally Lipschitz in t, which is equivalent to demonstrating that, for each x^a, t_0, and each $\delta \in (0, 1)$,

$$J^\star(x^a, t_0 + \delta) - J^\star(x^a, t_0) = O(\delta)$$

As a consequence of Bellman's principle of optimality, for any input u, we have,

$$\begin{aligned} J^\star(x^a, t_0) \le \int_{t_0}^{t_0+\delta} & c(x^a(t), u(t), t)\, dt \\ & + J^\star(x^a(t_0 + \delta), t_0 + \delta) \end{aligned} \tag{33}$$

For any initial state $(x^a(t_0))$, it follows from Lemma A.1 that there is a C^∞ input u such that for all $t \in [t_0, t_0 + \delta]$, the input is *cheap*, i.e., $u_\sigma(t) = 0$, and the state $(x^a(t))$ satisfies the uniform bound (32), with $x(t_0 + \delta) = x + O(\delta)$, $z(t_0 + \delta) = z$. For this input u, the integral term in (33) is of order $O(\delta)$, and since $J^\star$ is locally Lipschitz in x, we obtain the bound $J^\star(x^a, t_0) \le J^\star(x^a, t_0 + \delta) + O(\delta)$. It remains to establish the reverse inequality:

$$J^\star(x^a, t_0 + \delta) \le J^\star(x^a, t_0) + O(\delta). \tag{34}$$

First, recall the time-dependency in the cost comes from $\ell(t)$:

$$c(x^a, u, t) = c_{\mathrm{x}}(x) + c_{\mathrm{g}}(\ell(t) - z_\sigma) + \tfrac{1}{2}\kappa(u_\sigma - \ell'(t))^2$$

Let $\{(x^\circ(t), z^\circ(t), u^\circ(t)) : t_0 \le t \le \mathcal{T}\}$ denote the optimal solution that minimizes the value function, so that

$$J^\star(x^a, t_0) = \int_{t_0}^{\mathcal{T}} c(x^\circ(t), z^\circ(t), u^\circ(t), t)\, dt \tag{35a}$$

$$= \int_{t_0}^{\mathcal{T}} c(x^\circ(t), z^\circ(t), u^\circ(t), t + \delta)\, dt + O(\delta) \tag{35b}$$

Equation (35b) follows from Lipschitz continuity of ℓ and ℓ' (and recall that $c(x^a, u, t) := 0$ for $t > \mathcal{T}$ by Assumption 1).

For any feasible state-input trajectory $\{(x^a(t), u(t)) : t_0 + \delta \le t \le \mathcal{T}\}$ such that $x(t_0 + \delta) = x$, $z(t_0 + \delta) = z$, it follows from the definition of $J^\star$ that:

$$J^\star(x^a, t_0 + \delta) \le \int_{t_0+\delta}^{\mathcal{T}} c(x^a(t), u(t), t)\, dt. \tag{36}$$

For $t_0 + \delta \le t \le \mathcal{T}$, set

$$(x(t), z(t), u(t)) := (x^\circ(t - \delta), z^\circ(t - \delta), u^\circ(t - \delta))\,,$$

where the right hand side is based on the optimal solution in (35a). This trajectory satisfies $x(t_0 + \delta) = x$ and $z(t_0 + \delta) = z$, and consequently,

$$J^\star(x^a, t_0 + \delta) \le \int_{t_0+\delta}^{\mathcal{T}} c(x^\circ(t - \delta), z^\circ(t - \delta), u^\circ(t - \delta), t)dt \tag{37a}$$

$$= \int_{t_0}^{\mathcal{T}-\delta} c(x^\circ(t), z^\circ(t), u^\circ(t), t + \delta)\, dt \tag{37b}$$

$$\le \int_{t_0}^{\mathcal{T}} c(x^\circ(t), z^\circ(t), u^\circ(t), t + \delta)\, dt \tag{37c}$$

$$= J^\star(x^a, t_0) + O(\delta). \tag{37d}$$

The first inequality (37a) follows from (36); a change of variables leads to (37b); the fact that the cost function is non-negative results in (37c); and (37d) is a consequence of (35b).

Equation (37) gives the desired inequality (34). ■

Proof of Prop. II.1. Part (i) of Prop. II.1 is contained in Lemma A.3, so it remains to prove. It is clear from the definitions that $K^\star(x^a_\sigma, t_0) \le J^\star(x^a, t_0)$ for each x^a, t_0; we establish next the reverse inequality. For this we apply Lemma A.1: take any pair (x^+, z^+) satisfying $x_\sigma(t_0) = x^+_\sigma$, $z_\sigma(t_0) = z^+_\sigma$, and let u denote the input described in Lemma A.1, so that in particular $x^\delta(t_0+\delta) = x^+ + O(\delta)$ and $z^\delta(t_0+\delta) = z^+$. We obtain via Bellman's principle of optimality,

$$\begin{aligned} J^\star(x^a, t_0) &\le \int_{t_0}^{t_0+\delta} c(x^\delta(t), z^\delta(t), u^\delta(t))\, dt \\ &\quad + J^\star(x^\delta(t_0+\delta), z^\delta(t_0+\delta), t_0+\delta) \\ &= \int_{t_0}^{t_0+\delta} c(x^\delta(t), z^\delta(t), u^\delta(t))\, dt \\ &\quad + J^\star(x^+ + O(\delta), z^+, t_0+\delta) \\ &\le \int_{t_0}^{t_0+\delta} c(x^\delta(t), z^\delta(t), u^\delta(t))\, dt \\ &\quad + J^\star(x^+, z^+, t_0) + O(\delta) \end{aligned} \tag{38}$$

where the final inequality follows from Lemma A.3, i.e., the fact that $J^\star$ is locally Lipschitz in all its variables.

Based on the form of the cost function (11a) and the uniform bound on x^δ and z^δ (independent of δ) as a consequence of Lemma A.1, we may let $\delta \downarrow 0$ to obtain

$$J^\star(x^a, t_0) \le J^\star(x^+, z^+, t_0)$$

Taking the infimum over all (x^+, z^+) leads to $J^\star(x^a, t_0) \le K^\star(x^a_\sigma, t_0)$ as claimed. ■

B. Co-state dynamics and state space collapse

The dynamics of the dual variables appearing in Prop. III.2 are given in the following lemma.

Lemma A.4. *Let $t_0 \in (0, \mathcal{T})$. The optimal input-state $(u^\star, x^\star, z^\star)$ and dual variables $\{\lambda^\star, \beta^\star\}$ satisfy the following co-state equations:*

$$\tfrac{d}{dt}\lambda^\star_i(t) = \alpha_i \lambda^\star_i(t) - c_i'(x^\star(t)), \tag{39a}$$

$$\tfrac{d}{dt}\beta^\star_i(t) = \lambda^\star_i(t) + c_g'(\ell(t) - z^\star_\sigma(t)), \quad t \in [t_0, \mathcal{T}], \tag{39b}$$

with boundary condition $\lambda^\star_i(\mathcal{T}) = 0,\ \beta^\star_i(\mathcal{T}) = 0$ for each i.

Proof. The state $(x^\star(t), z^\star(t))$ is continuously differentiable and the control $u^\star(t)$ is continuous on $[t_0, \mathcal{T}]$ as a consequence of Assumption 2. In addition, the dynamics, eqs. (9d) and (9e), are linear and hence continuously differentiable with respect to each of the variables. Further, Assumption 1 implies that the cost function $\boldsymbol{c}$ in (11a) is continuously differentiable with respect to all the variables. Consequently, the optimal input-state pair $u^\star, (x^\star, z^\star)$ satisfies Pontryagin's minimum principle on the closed interval $[t_0, \mathcal{T}]$ [26, Section 4.2]. The rest of the proof follows from this result. In particular, the minimum principle implies

$$\begin{aligned} \tfrac{d}{dt}\lambda^\star_i(t) &= -\tfrac{\partial}{\partial x} H\big(x^\star(t), z^\star(t), u^\star(t), \lambda^\star(t), \beta^\star(t), t)\big), \\ \tfrac{d}{dt}\beta^\star_i(t) &= -\tfrac{\partial}{\partial z} H\big(x^\star(t), z^\star(t), u^\star(t), \lambda^\star(t), \beta^\star(t), t)\big) \end{aligned}$$

which yields (39a). The boundary conditions $\lambda^\star_i(\mathcal{T}) = \beta^\star_i(\mathcal{T}) = 0$ hold because there is no terminal cost [43, Theorem 1]. ■

The optimal input is considered next.

Lemma A.5. *Let $t_0 \in (0, \mathcal{T})$. The optimal input on $[t_0, \mathcal{T}]$ is*

$$u^\star_\sigma(t) = \tfrac{d}{dt}\ell(t) - \frac{1}{\kappa}\beta^\star_i(t), \qquad t \in [t_0, \mathcal{T}] \tag{40}$$

Moreover, $\beta^\star(t) := \beta^\star_i(t)$, $\lambda^\star(t) := \lambda^\star_i(t)$ are independent of i.

Proof. Similar to the proof of Lemma A.4, the conclusion (40) is a consequence of Pontryagin's minimum principle, and the first-order condition for minimality:

$$0 = \tfrac{\partial}{\partial u} H\big(x^\star(t), z^\star(t), u^\star(t), \lambda^\star(t), \beta^\star(t), t)\big)$$

This establishes (40), which then implies that $\beta^\star_i(t)$ is independent of i for each t. The conclusion that $\lambda^\star_i(t)$ is also independent of i follows from eq. (39b). ■

The lemma reinforces the co-state collapse identified in Prop. III.2. Proof of the main result is provided next.

Proof of Thm. III.1. Equations (16a) and (16b) are the state equations. Equations (16c) and (16d) and the final time boundary conditions on the co-state variables follow from Lemmas A.4 and A.5. The optimal feedback policy (16e) is obtained from Lemma A.5.

As a consequence of Assumption 2 and the fact that the cost functional is convex with respect to the control and that the dynamics are linear, the solution $(x^\star, z^\star, \lambda^\star, \beta^\star, u^\star)$ satisfying Pontryagin's minimum principle is both necessary and sufficient for optimality [6, Chapter 7]. ■

C. Quadratic cost functions:

Representations of the value function, and the optimal co-state/input are provided next, in the setting of Section IV-A:

Lemma A.6. *Consider the cost function $c(x(t), z(t), u(t), t)$ provided in* (27) *and the value function J^* defined in* (26) *with $\ell(t) = \overline{\ell}$ for all $t \in [0, \mathcal{T}]$. Then, for $t \in (0, \mathcal{T}]$, we have the following:*

(i) *The value function has the form:*

$$J^\star(x^a, t) = K^\star(x^a_\sigma, t) = \tfrac{1}{2}\begin{bmatrix} x_\sigma \\ z_\sigma \end{bmatrix}^{\intercal} P_t \begin{bmatrix} x_\sigma \\ z_\sigma \end{bmatrix}$$

where P_t is a symmetric, positive semi-definite matrix.

(ii) *The optimal co-state $(\lambda^\star(t), \beta^\star(t))$ can be represented in terms of the optimal state trajectories:*

$$\begin{bmatrix} \lambda^\star(t) \\ \beta^\star(t) \end{bmatrix} = P_t \begin{bmatrix} x^\star_\sigma(t) \\ z^\star_\sigma(t) \end{bmatrix} = P_t w^{\intercal} \begin{bmatrix} x^\star(t) \\ z^\star(t) \end{bmatrix} \tag{41}$$

(iii) *The optimal input $u^\star_\sigma(t)$ is*

$$u^\star_\sigma(t) = -\frac{1}{\kappa M}\mathbf{1}^{\intercal} B^{\intercal} w P_t w^{\intercal} \begin{bmatrix} x^\star(t) \\ z^\star(t) \end{bmatrix} \tag{42}$$

Proof. Statement (i) follows from the fact that for a linear system with quadratic costs, the value function has a quadratic form [26, Section 6.1.3]. Then, in the second result (41), the first equality follows from the definitions of $\lambda^\star$ and $\beta^\star$ in (14), and the second equality follows from (2).

From basic matrix algebra and (41), it follows that the right-hand side of (42) equals $-\beta^\star(t)/\kappa$. Since $\frac{d}{dt}\ell(t) = 0$ by assumption, (16e) leads to the third result (42). ■

Proof of Prop. IV.1. Statement (i) of Lemma A.6 yields (28). From (16c) and (41), we can write:

$$w\tfrac{d}{dt}\begin{bmatrix}\lambda^\star(t)\\ \beta^\star(t)\end{bmatrix} = \begin{bmatrix}-Q - A^\intercal w P_t w^\intercal\end{bmatrix}\begin{bmatrix}x^\star(t)\\ z^\star(t)\end{bmatrix} \tag{43}$$

Differentiating (41) with respect to time, it follows from (3), (42), and matrix algebra that,

$$\begin{aligned} w\tfrac{d}{dt}\begin{bmatrix}\lambda^\star(t)\\ \beta^\star(t)\end{bmatrix} = \Big[&w\dot{P}_t w^\intercal + wP_t w^\intercal A \\ &- \frac{1}{\kappa} w P_t \begin{bmatrix}0 & 0\\ 0 & 1\end{bmatrix} P_t w^\intercal\Big]\begin{bmatrix}x^\star(t)\\ z^\star(t)\end{bmatrix}\end{aligned} \tag{44}$$

Denote $W := w^\intercal w$ (an invertible matrix). The differential Riccati equation (29) follows from (43) and (44) after multiplication by $W^{-1}w^\intercal$ on the left and wW^{-1} on the right.

The boundary condition $P_{\mathcal{T}} = 0$ follows from the fact that $K^\star(x^a_\sigma, \mathcal{T}) = 0$ for any x, z ■

Proof of Prop. IV.2. From (17) and (41),

$$\begin{aligned} x^\star(t) = \Big[&-Q_x^{-1}A_x[\mathbf{1}\quad \mathbf{0}]P_t - Q_x^{-1}[\mathbf{1}\quad \mathbf{0}]\dot{P}_t\Big]\begin{bmatrix}x^\star_\sigma(t)\\ z^\star_\sigma(t)\end{bmatrix} \\ &- Q_x^{-1}[\mathbf{1}\quad \mathbf{0}]P_t\tfrac{d}{dt}\begin{bmatrix}x^\star_\sigma(t)\\ z^\star_\sigma(t)\end{bmatrix},\end{aligned} \tag{45}$$

where $Q_x = \mathrm{diag}(q_i)$ and $A_x = -\mathrm{diag}(\alpha_i)$. Rewriting (8a) and (8b) in matrix-vector form gives

$$\tfrac{d}{dt}\begin{bmatrix}x^\star_\sigma(t)\\ z^\star_\sigma(t)\end{bmatrix} = \begin{bmatrix}\mathbf{1}^\intercal A_x x^\star(t) - z^\star_\sigma(t)\\ u^\star_\sigma(t)\end{bmatrix} \tag{46}$$

Substituting $x^\star(t)$ from (45) and $u^\star_\sigma(t)$ from (42) in (46) yields the first result (30) after simplification. Since (30) is a linear time-varying system, there exists a unique state-transition matrix $\phi(t,0)$ such that,

$$\begin{bmatrix}x^\star_\sigma(t)\\ z^\star_\sigma(t)\end{bmatrix} = \phi(t,0)\begin{bmatrix}x^\star_\sigma(0)\\ z^\star_\sigma(0)\end{bmatrix} \tag{47}$$

The map G^1 in (31) follows from (45), and G^2 follows from (16a). ■

REFERENCES

[1] California ISO – Supply, Today's Outlook. https://www.caiso.com/TodaysOutlook/Pages/supply.html. Accessed: 2022-10-14.

[2] H. Ballouz, J. Mathias, S. Meyn, R. Moye, and J. Warrington. Addressing misconceptions on the performance of the energy market in Texas. Utility Dive: https://www.utilitydive.com/news/addressing-misconceptions-on-the-performance-of-the-energy-market-in-texas/598436/, April 15 2021.

[3] H. Ballouz, J. Mathias, S. Meyn, R. Moye, and J. Warrington. Reliable power grid: Long overdue alternatives to surge pricing. *arXiv 2103.06355*, March 2021.

[4] P. Barooah, A. Bušić, and S. Meyn. Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid. In *Proc. 48th Annual Hawaii International Conference on System Sciences (HICSS)*, pages 2700–2709, Kauai, Hawaii, 2015.

[5] E. Benenati, M. Colombino, and E. Dall'Anese. A tractable formulation for multi-period linearized optimal power flow in presence of thermostatically controlled loads. In *IEEE Conference on Decision and Control*, pages 4189–4194. IEEE, 2019.

[6] A. Bressan and B. Piccoli. *Introduction to the mathematical theory of control*, volume 1. American institute of mathematical sciences Springfield, 2007.

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

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

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

[10] 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.

[11] 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.

[12] I.-K. Cho and S. P. Meyn. Efficiency and marginal cost pricing in dynamic competitive markets with friction. *Theoretical Economics*, 5(2):215–239, 2010.

[13] F. H. Clarke and R. B. Vinter. The relationship between the maximum principle and dynamic programming. *SIAM Journal on Control and Optimization*, 25(5):1291–1311, 1987.

[14] A. R. Coffman, N. Cammardella, P. Barooah, and S. Meyn. Aggregate flexibility capacity of TCLs with cycling constraints. *IEEE Transactions on Power Systems*, 38(1):52–62, 2023.

[15] F. E. R. Commission et al. FERC opens wholesale markets to distributed resources: Landmark action breaks down barriers to emerging technologies, boosts competition. *FERC News Release*, September 17, 2020.

[16] J. Engel. SPP sets grid operator record with 90% renewable energy penetration. https://www.renewableenergyworld.com/wind-power/spp-sets-grid-operator-record-with-90-renewable-energy-penetration/#gref, March 2022.

[17] L. A. D. Espinosa and M. Almassalkhi. A packetized energy management macromodel with quality of service guarantees for demand-side resources. *IEEE Transactions on Power Systems*, 35(5):3660–3670, 2020.

[18] L. A. D. Espinosa, A. Khurram, and M. Almassalkhi. Reference-tracking control policies for packetized coordination of heterogeneous der populations. *IEEE Transactions on Control Systems Technology*, 29(6):2427–2443, 2021.

[19] A. Faruqui, R. Hledik, and J. Tsoukalis. The power of dynamic pricing. *The Electricity Journal*, 22(3):42–56, 2009.

[20] B. Francis. The optimal linear-quadratic time-invariant regulator with cheap control. *IEEE Trans. Automat. Control*, 24(4):616–621, 1979.

[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.

[22] M. L. Hautus and L. M. Silverman. System structure and singular control. *Linear algebra and its applications*, 50:369–402, 1983.

[23] A. Jameson and R. O'Malley. Cheap control of the time-invariant regulator. *Applied Mathematics and Optimization*, 1(4):337–354, 1975.

[24] I. Karatzas. A class of singular stochastic control problems. *Advances in Applied Probability*, 15(2):225–254, 1983.

[25] N. Kumar, P. Besuner, S. Lefton, D. Agan, and D. Hilleman. Power plant cycling costs. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2012.

[26] D. Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2011.

[27] H. Lo, S. Blumsack, P. Hines, and S. Meyn. Electricity rates for the zero marginal cost grid. *The Electricity Journal*, pages 39–43, 2019.

[28] F. Lu, J. Mathias, S. Meyn, and K. Kalsi. Model-free characterizations of the Hamilton-Jacobi-Bellman equation and convex Q-learning in continuous time. *arXiv.2210.08131*, 2022.

[29] R. Malhame and C.-Y. Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9):854–860, Sep. 1985.

[30] J. Mathias. *Balancing the power grid with distributed control of flexible loads*. PhD thesis, University of Florida, Gainesville, FL, USA, 2022.

[31] J. Mathias, A. Bušić, and S. Meyn. Load-level control design for demand dispatch with heterogeneous flexible loads. *IEEE Transactions on Control Systems Technology*, pages 1–14, 2023.
[32] J. Mathias, R. Kaddah, A. Bušić, and S. Meyn. Smart fridge / dumb grid? Demand Dispatch for the power grid of 2020. In *Proc. 49th Annual Hawaii International Conference on System Sciences (HICSS)*, pages 2498–2507, Jan 2016.
[33] J. Mathias, R. Moye, S. Meyn, and J. Warrington. State space collapse in resource allocation for demand dispatch. In *Proc. of the Conf. on Dec. and Control*, pages 6181–6188 (and arXiv:1909.06869), Dec 2019.
[34] J. Mathieu, S. Koch, and D. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Trans. Power Systems*, 28(1):430–440, 2013.
[35] J. L. Mathieu, M. E. Dyson, and D. S. Callaway. Resource and revenue potential of California residential load participation in ancillary services. *Energy Policy*, 80(0):76 – 87, 2015.
[36] 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.
[37] S. Meyn, T. Samad, and J. Stoustrup, editors. *IMA volume on the control of energy markets and grids*. Springer, 2018.
[38] S. P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, 2007. Pre-publication edition available online.
[39] M. Negrete-Pincetic, G. Wang, M. Arancibia, A. Kowli, E. Shafieepoorfard, and S. Meyn. The value of volatile resources in electricity markets. *Sustainable Energy, Grids and Networks*, 11:46 – 57, 2017.
[40] M. I. Reiman. Some diffusion approximations with state space collapse. In F. Baccelli and G. Fayolle, editors, *Modelling and Performance Evaluation Methodology*, pages 207–240, Berlin, Heidelberg, 1984. Springer Berlin Heidelberg.
[41] A. W. Roberts and D. E. Varberg. Another proof that convex functions are locally lipschitz. *The American Mathematical Monthly*, 81(9):1014–1016, 1974.
[42] V. Saksena, J. O'Reilly, and P. Kokotovic. Singular perturbations and time-scale methods in control theory: Survey 1976–1983. *Automatica*, 20(3):273 – 293, 1984.
[43] R. B. Vinter. Optimal control and Pontryagin's Maximum Principle. *Encyclopedia of Systems and Control*, pages 950–956, 2015.
[44] J. Willems, A. Kitapci, and L. Silverman. Singular optimal control: a geometric approach. *SIAM Journal on Control and Optimization*, 24(2):323–337, 1986.

Joel Mathias is currently a postdoctoral research scholar at Arizona State University. He received a Ph.D. degree in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in Spring 2022. He also has an M.S. degree in electrical and computer engineering from the University of Florida and B.Eng. degree in electronics and communications engineering from the University of Mumbai.

He was a research associate with the Tata Institute of Fundamental Research, Mumbai, India, in 2011–2012, and his industry experience includes Electric Power Engineers, Inc., Austin, TX, USA, and Tata Consultancy Services, Mumbai, India. For his graduate studies in the USA, he has received the J.N. Tata Endowment for the higher education of Indians, the Lady Navajbai Tata scholarship, and achievement awards from the University of Florida.

His research interests are in stochastic processes, control theory, and optimization, with applications in the area of smart power grids.

Sean Meyn was raised by the beach in Santa Monica, California. Following his BA in mathematics at UCLA, he moved on to pursue a PhD with Peter Caines at McGill University. After about 20 years as a professor of ECE at the University of Illinois, in 2012 he moved to beautiful Gainesville. He is now Professor and Robert C. Pittman Eminent Scholar Chair in the Department of Electrical and Computer Engineering at the University of Florida, director of the Laboratory for Cognition and Control, and Inria International Chair at Inria, France. He is an IEEE CSS distinguished lecturer. His interests span many aspects of stochastic control, stochastic processes, information theory, and optimization. For the past decade, his applied research has focused on engineering, markets, and policy in energy systems.

Robert Moye has been in the energy industry for over forty years, getting his start with utilities in Florida where he focused on long-term resource planning. Since the energy markets were re-regulated in 1994, he has worked for electricity and natural gas energy marketing companies that manage assets across the United States, Canada, and Mexico. Bob is currently the Senior Vice President of Energy Management for Tyr Energy, Inc. In this role he oversees the management of thermal and renewable energy assets for electric utilities, independent power producers and industrial facilities in North America. Bob earned a BS and ME in Electrical Engineering, an MBA, and a PhD in Electrical Engineering from the University of Florida. Bob is also a registered Professional Engineer in the State of Florida.

Joseph Warrington was until December 2019 a Senior Scientist in the Automatic Control Lab (IfA) at ETH Zurich, Switzerland, and is now Principal Optimisation Engineer at AstraZeneca, Cambridge, UK. His Ph.D. is from ETH Zurich (2013), and his B.A. and M.Eng. degrees in Mechanical Engineering are from the University of Cambridge (2008). From 2014-2016 he worked as an energy consultant at Baringa Partners LLP, London, UK, and he has also worked as a control systems engineer at Wind Technologies Ltd., Cambridge, UK, and privately as an operations research consultant. He is the recipient of the 2015 ABB Research Prize for an outstanding PhD thesis in automation and control, and a Simons-Berkeley Fellowship for the period January-May 2018. His research interests include dynamic programming, large-scale optimization, and predictive control, with applications including power systems and transportation networks.