



HAL
open science

Distributed Control Design for Balancing the Grid Using Flexible Loads

Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Bušić, Sean Meyn

► **To cite this version:**

Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Bušić, Sean Meyn. Distributed Control Design for Balancing the Grid Using Flexible Loads. S. Meyn; T. Samad; I. Hiskens; J. Stoustrup. IMA Volume on the Control of Energy Markets and Grids, 162, Springer, pp.383-411, 2018. hal-01656726

HAL Id: hal-01656726

<https://hal.science/hal-01656726>

Submitted on 6 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed Control Design for Balancing the Grid Using Flexible Loads

Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Bušić, and Sean Meyn

Abstract Inexpensive energy from the wind and the sun comes with unwanted volatility, such as ramps with the setting sun or a gust of wind. Controllable generators manage supply-demand balance of power today, but this is becoming increasingly costly with increasing penetration of renewable energy. It has been argued since the 1980s that consumers should be put in the loop: “demand response” will help to create needed supply-demand balance. However, consumers use power for a reason, and expect that the quality of service (QoS) they receive will lie within reasonable bounds. Moreover, the behavior of some consumers is unpredictable, while the grid operator requires predictable controllable resources to maintain reliability. The goal of this chapter is to describe an emerging science for *demand dispatch* that will create *virtual energy storage* from flexible loads. By design, the grid-level services from flexible loads will be as controllable and predictable as a generator or fleet of batteries. Strict bounds on QoS will be maintained in all cases. The potential economic impact of these new resources is enormous. California plans to spend billions of dollars on batteries that will provide only a small fraction of the balancing services that can be obtained using demand dispatch. The potential impact on society is enormous: a sustainable energy future is possible with the right mix of infrastructure and control systems.

Acknowledgements Research supported by the DOE Buildings Technologies Office as part of the Virtual Batteries project under the Buildings-to-grid Transactive Energy program (the authors would like to thank Mr. Joe Hagerman for his continued support); the National Science Foundation under grants CPS-0931416 and EPCN-1609131, the French National Research Agency grant ANR-16-CE05-0008, and by PGM0. Our thanks to the anonymous reviewers for their suggestions, and to Robert Moyer at UF and Rainbow Energy for his encouragement and commentary.

Yue Chen, Joel Mathias, and Sean Meyn
Department of ECE, University of Florida, Gainesville

Md Umar Hashmi and Ana Bušić
Inria Paris
DI ENS, École normale supérieure, CNRS, PSL Research University, Paris

1 Introduction

As more wind and solar energy come online, the system operators who run the power grid are faced with a problem: how do they compensate for the variable nature of the sun and wind?

Low frequency variability from solar gives rise to the famous “duck curve” anticipated at CAISO [8]: It is predicted that ramps of 30% of the load over a few hours may be commonplace. In 2011, they introduced new market rules for pricing *flexible ramping products* to help combat this volatility. Fig. 1 is taken from a 2014 presentation at the Southwest Power Pool (SPP) working group meeting.¹ The net-load on March 8, 2014 shows the emergence of the CAISO duck-curve. The price data illustrates that insufficient ramping resources can cause enormous spikes in wholesale power prices.




Fig. 1 Interaction between power & price dynamics at CAISO.

MISO (an ISO in central North America) recently alerted FERC to the need for new market rules to incentivize ramping products. They argue that the need for these services is increasing with the introduction of renewable energy: “*under its current market structure, short-term Net Load variations could create a situation where dispatchable resources have sufficient capacity, but there is a short-term scarcity event because MISO has inadequate ramp capability to respond to unexpected variations in Net Load ... such ramp capability shortages could result in ... dispatch intervals during which the price of energy can increase significantly due to scarcity pricing, even if the event does not present a significant reliability risk*”². MISO is not concerned about energy: they are lacking responsive resources, even though there is sufficient capacity.

The control systems diagram in Fig. 2 provides a simple view of how the grid is operated today, in which wind and solar are viewed as sources of disturbances. In North America, the **GRID** block is in fact a subset of the grid called a *balancing region*. The block denoted **Compensation** represents a *balancing authority* (BA).

¹ From Tony Delacluyse of PCI, with permission.

² <http://tinyurl.com/FERC-ER14-2156-000>

The grid-level measurements obtained by the BA are summarized as a scalar function of time called the area control error (ACE). It is a linear combination of two error signals: the deviation of local grid frequency from the nominal 60 Hz, and the tie-line error — defined as the mismatch between scheduled and actual flow of power out of the balancing region. Command signals are broadcast to resources such as controllable generators so that the ACE signal is kept within desired bounds.

The compensator G_c is designed by the BA in a particular region. For example, PJM (an RTO in the Eastern U.S.) creates their RegA and RegD signals by passing the ACE signal first through a PI compensator, and then through a bandpass filter. In this case the compensator G_c in Fig. 2 is taken to be a PI controller, and the bandpass filters are embedded in the block denoted **Actuation**. The decomposition “ $H = H_a + H_b + \dots$ ” represents many resources acting in parallel to provide actuation, which may include controllable generation and batteries.

It is anticipated that the basic architecture illustrated in Fig. 2 will remain in place for many decades to come. The grid will become more adaptable to persistent disturbances or crisis through a combination of control techniques and hardware.




Fig. 2 Power Grid Control Loop. A question addressed in this chapter: *where do we find H?*

The term *ancillary service* refers to resources required to maintain supply-demand balance in the grid, but do not necessarily supply energy. While controllable generation is the most common source of most ancillary services today, other technologies such as flywheels and batteries are increasingly popular because of their performance, and because of new state and federal policies. The FERC report [37, pp. 23-24] contains a survey of experiments conducted by Beacon Power and Primus Power on the value of highly responsive resources for ancillary service. Primus claims approximately 76 percent more ACE correction compared with what can be obtained from generation sources. This is because a slower “... *resource lags to the point of working against needed ACE correction.*” In addition to their poorer performance, the use of generation for ancillary service may be costly in terms of fuel and emissions because the generators are not running at their ideal conditions, and in fact may be online only to help regulate the grid.

The California Public Utilities Commission has encouraged utilities to invest billions of dollars on batteries, and many have responded. Southern California Edison (SCE) recently announced the “largest battery in the world.” It is capable of storing 32 MWhs (megawatt-hours) of energy. It is not a single battery, but a 6,300 square-foot facility that houses an enormous array of lithium-ion batteries.

Purpose of this chapter. Batteries may be a clean source of ancillary service, but currently they are an expensive solution. In addition to the large space required for large systems, batteries have finite life time, and waste energy as they are charged and discharged to service the grid [17].

Distributed control architectures are described in this chapter to create *virtual energy storage* (VES) based on the inherent flexibility in power consumption from a majority of loads. The ancillary services that can be obtained include regulation (such as automatic generation control, or AGC), smooth peaks in load, address ramps from wind or solar generation, and help to recover gracefully from contingencies such as generation faults. It is believed that VES is a low-cost complement to batteries and power plants, and may in the future provide the majority of required ancillary services.

The term *Demand Dispatch* is a convenient alternative to *Demand Response*; the latter is defined by policy makers and regulatory bodies (such as FERC) as load-shedding in exchange for some monetary reward. *Load shedding is not the goal of the technology surveyed here.* In applications to both regulation and ramping services, the distributed control algorithms are designed so that power consumption is increased and decreased over time, while keeping the total energy deviation over time at zero — *just like charging and discharging of a battery.*

The control architectures described in this chapter are based on a series of papers on distributed control [35, 36, 13, 11]. The proposed frequency decomposition of VES resources was first introduced in [20, 19] in the context of commercial buildings, and generalized in [1]. The key novel contribution in all of this work is the focus on “intelligence at the load”, manifested by local control loops. There are many benefits:

- (i) Communication infrastructure requirements are reduced, which hopefully leads to both improved security, and higher consumer confidence regarding privacy.
- (ii) A simple control problem at the BA, since the single-input / single-output system is highly controllable.
- (iii) Strict bounds on quality of service (QoS) to the consumer are guaranteed.

This chapter does not consider market issues. It is assumed that consumer engagement will be achieved through contractual agreements and periodic credits, such as those offered by Florida Power and Light in their OnCall[®] program.

Some history and further reading. In the early eighties, Schweppe wrote a series of influential articles on the value of new architectures for the grid [41, 40], with an emphasis on demand response based on either automation or prices. Tools for analysis were lacking at the time, but many researchers came to fill the void. An influential example is the paper [29], that introduced ideas from statistical mechanics to model a large population of thermostatically controlled loads (TCLs).

There is substantial literature on indirect load control, where customers are encouraged to shift their electricity usage in response to real-time prices (several papers in this volume survey this literature, including the articles authored by Spence and by Moye). Dynamic prices can introduce uncertain dynamics, such as cyclical

price fluctuations and increased sensitivity to exogenous disturbances, and present a risk to system stability [39, 9, 15].

Randomization is an essential element of the distributed control architecture described in this chapter. Its value has been widely recognized in academia as well as industry [42].

On the academic side, Matheiu’s dissertation [32] and references [33, 34] were highly influential, motivating in part the research surveyed in this chapter and others [16, 26, 46, 48]. The control model in [32] is based on the mean-field setting of [29], with the introduction of a control signal from a central authority: at each time slot, a BA or aggregator broadcasts probability values $\{p_\tau^\oplus, p_\tau^\ominus : \tau \in \mathbb{R}\}$ where p_τ^\oplus (p_τ^\ominus) denotes the probability of turning the device on (off) when the temperature of the device is τ . The temperatures are binned to obtain a finite state-space aggregate model. This model is bilinear and partially observed, where the state x is the histogram of load temperature and power consumption. The bilinear control system is transformed to a linear model by defining products of probability and state as an input. The resulting linear state space model has the same state, but the vector-valued input is now defined as products of the form $u_k = p_\tau^m x_j$ for some $\tau(k)$, $j(k)$, and $m(k) \in \{\oplus, \ominus\}$. Feedback control design is performed based on LQR. However, it is still necessary to recover the probability vector $\{p_\tau^m\}$. In this prior work, this is defined as the ratio of components of the input $u(t)$, and components of the *estimate* of the state at time t (see e.g. eq. (11) of [34]). It is assumed that estimates are computed by the BA based on measurements of aggregate power consumption. A current challenge with this approach is the creation of sufficiently accurate state estimates for an inherently infinite-dimensional system. Challenges to state estimation are discussed in [13], where it is shown that the linearized mean field model may not be observable. Robustness of this approach to bilinear control systems is another important area for future research.

The approach to distributed control surveyed in Sections 2 and 3 involves an entirely different approach to local control at each load. One example is the *Individual Perspective Design* (IPD) described in Section 3.2. This can be regarded as an application of the MDP technique of Todorov [44], but only in one special case: the construction of [44] can be applied only if there is no exogenous stochastic disturbance in the load model. Contained in Section 3.2 are techniques to extend this design to a broader class of load models. These ideas were first applied to demand-dispatch in [35], and have seen many extensions since. For more history the reader is referred to [36, 14], in addition to the papers surveyed in Section 3.2. While beyond the scope of this article, it is important to note that Todorov’s ‘linearly solvable’ MDP model [44] is similar to prior work such as [24], and the form of the solution could have been anticipated from well-known results in the theory of large-deviations for Markov chains [7]. It is pointed out in [45] that this approach has a long history in the context of controlled stochastic differential equations [18].

The remainder of the chapter consists of six sections organized as follows. Section 2 contains a high-level description of the control architecture, with details on distributed control contained in Section 3. The next three sections contain examples of distributed control of a large collection of resources in three settings: Section 4

concerns TCLs, Section 5 presents application to a population of residential pool pumps, and Section 6 describes application of similar methodology to a spatially distributed population of batteries. Conclusions and discussion of future research is contained in Section 7.

2 Distributed Control Architecture

This section contains a list of specific goals, and general control design techniques that offer solutions.

Problem statement. The grid operator requires resources to balance the grid at all times. The hypothesis of this chapter is that a large proportion of the needed resources can come in the form of virtual storage from flexible loads. Reliable grid services can be obtained from loads, but this requires a well-designed control architecture.

A particular hierarchical control architecture is developed in this chapter. One realization is illustrated by the feedback structure shown in Fig. 2, in which the actuation block is composed of many resources acting in parallel, including generation, batteries, and virtual energy storage.

Assumptions regarding this control structure include

- (i) *Local control:* This will be based in part on randomized decision rules. This provides necessary degrees of freedom in shaping aggregate dynamics. Randomization also helps to prevent synchronization of the response from loads.
- (ii) *Information flow from loads:* Two-way information exchange between the BA and individual loads is *not* a component of this architecture. In [36] it is assumed that the BA measures aggregate power consumption from the loads under its authority. Alternatively, each load broadcasts its power state several times per day, and aggregate power consumption is estimated at the BA [13].

In [30, 31] it is argued that it is possible to create a reliable control system in which direct information flow from loads to BA is entirely absent. This requires more complex control at each load, and hence is beyond the scope of this chapter.

- (iii) *Information flow from the BA:* A single regulation signal is broadcast to each collection of loads of the same class, as illustrated in Fig. 3. This signal is designed based on grid level measurements, and a model of the aggregate behavior of the loads in each class.


The value of “local intelligence” at each load is vital for the envisioned architecture. Feedback loops at each load are used to ensure that QoS constraints are met, and also so that the aggregate of loads will appear to the grid operator as a reliable resource – much like a battery system, or a controllable generator.

Consumer choice will be an input to any VES system, and a monetary reward may be part of the arrangement. A contract for services can be established so that

the consumer is rewarded for participation, without exposing him or her to the complexity and uncertainty of the grid. In this way the BA or aggregator can design the system so that highly reliable grid services are obtained, while respecting the QoS constraints of the consumer.

In the future it is possible that some loads will be grid-friendly by design; the consumer will never know that their refrigerator is helping to regulate the grid.

Fig. 3 *Control architecture*: a common command signal is transmitted to each load in a particular class. The resulting input-output system from ζ to power consumption y is regarded as virtual energy storage.



Local control The lowest level of control in the proposed architecture is at an individual load, such as a water heater, refrigerator, agricultural water pump, or air-conditioner. The load is equipped with sensors. For example, the microprocessor in a water heater receives measurements of water temperature at one or more locations in the unit. It is also assumed that the load receives measurements from the grid. This could be purely local, such as the grid frequency measured locally [30, 31]. Theory is best developed in the setting where each load receives a signal from the BA.

The local control loop is designed to meet these potentially conflicting goals: 1. Ensure that the load is providing the desired services to the consumer, respecting strict bounds on QoS, and 2. Ensure that the *aggregate* of loads responds to a signal from the BA in a manner that is both predictable and beneficial to the grid.

One obvious challenge: the degrees of freedom are extremely limited for a typical load of interest. For example, a residential water heater or refrigerator can be in only one of a small number of power states. Contained in Section 3 are several design techniques for local control that result in smooth aggregate behavior. This is possible without the introduction of complex scheduling rules, or the solution of real-time optimization problems at the BA.

Macro control This high-level control layer may be a part of the traditional BA, or through a load aggregator. The balancing challenges are of many different categories, on many different time-scales:

- (i) Automatic Generation Control (AGC): time scales of seconds to 20 minutes.
- (ii) Balancing reserves. In the Bonneville Power Authority, the balancing reserves include both AGC and balancing on timescales of many hours.³
- (iii) Contingencies (e.g., a generator outage)

The final two challenges are observed in Fig. 1:

³ Balancing on a slower time-scale is achieved through real time markets in some other regions of the U.S., and in every region under the jurisdiction of an RTO.

- (iv) Peak shaving.
- (v) Smoothing ramps from solar or wind generation.

In this chapter it is assumed that these high-level control problems are addressed as they are today: the BA receives measurements of the grid, and based on this information sends out signals to each resource in its domain. In many cases control loops are based on standard PI (proportional-integral) control design.

The difference here is that some resources are virtual, such as a collection of water heaters. A large collection of batteries distributed across the region might be regarded as a single resource – in this case, local control loops will be installed in each battery system so that the aggregate behaves as a single massive battery.

3 Mean-Field Control Design

Standard approaches for solving a stochastic control problem include stochastic dynamic programming and Markov decision processes (MDP) [38]. The future power grids will contain millions of smart components, which prohibits centralized decision making using these techniques as they do not scale well with the number of different components in the system (both the state space and the control space of the model grow exponentially with the number of components). The extension of MDP models to the case of optimization problems involving many agents that are making decisions based on partial knowledge of the system is called DEC-POMDP (decentralized partially observable MDP). These problems are NEXP-hard for the finite horizon optimization case [2], and undecidable in the infinite horizon case [28].

In physics and probability theory, mean field theory (MFT) approximates the behavior of a large number of small individual components which interact with each other. The effect of all the other individuals on any given individual is approximated by a single averaged effect, thus reducing a many-body problem to a one-body problem. The mean-field ideas first appeared in physics in the work of Pierre Curie and Pierre Weiss to describe phase transitions [23, 47]. Approaches inspired by these ideas have seen applications in epidemic models [3], computer network performance and game theory [27, 22]. In power systems, they were first used to model the aggregate dynamic of the collection of water-heaters in [29], and more recently in [25, 46, 43]. However, the global objective optimization under mean-field interactions remains very challenging and an exact analysis is possible only under restrictive assumptions on the local dynamics and the cost structure. There is still a significant gap between the theoretical assumptions and the applications, and the results may be sensitive to the modeling errors.

The approach in [11, 13, 6] combines the mean-field theory with classical feedback control. The main idea consists in defining a parametrized family of randomized local decision rules that lead to an aggregate behavior with desirable control properties (e.g. passivity for the linearized aggregate input-output system).

This section provides an overview of key concepts and results of this approach.

3.1 Mean-field model

A nominal Markovian model for an individual load is created based on its typical operating behavior. This is described as a Markov chain with transition matrix denoted P_0 , with state space $\mathsf{X} = \{x^1, \dots, x^d\}$. For example, a water chiller turns on or off depending upon the temperature of the water. In this case, a state value x^i encodes water temperature as well as the power state (on or off).

A family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$ is then constructed to define local decision making. Each load evolves as a controlled Markov chain on X , with common input $\boldsymbol{\zeta} = (\zeta_0, \zeta_1, \dots)$. It is assumed that the scalar signal $\boldsymbol{\zeta}$ is broadcast to each load. If a load is in state x at time t , and the value ζ_t is broadcast, then the load transitions to the state x' with probability $P_{\zeta_t}(x, x')$. Letting X_t^i denote the state of the i th load at time t , and assuming N loads, the empirical pmf (probability mass function) is defined as the average,

$$\mu_t^N(x) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{X_t^i = x\}, \quad x \in \mathsf{X}.$$

The mean-field model is the deterministic system defined by the evolution equations,

$$\mu_{t+1} = \mu_t P_{\zeta_t}, \quad t \geq 0, \quad (1)$$

in which μ_t is a row vector of dimension d . Under general conditions on the model and on μ_0 it can be shown that μ_t^N is approximated by μ_t .

In [11, 13, 36] it is assumed that average power consumption is obtained through measurements or state estimation: Let $\mathcal{U}(x)$ denote power consumption when the load is in state x , where $\mathcal{U} : \mathsf{X} \rightarrow \mathbb{R}_+$. The average power consumption is denoted

$$y_t^N = \frac{1}{N} \sum_{i=1}^N \mathcal{U}(X_t^i),$$

which is approximated using the mean-field model:

$$y_t = \sum_x \mu_t(x) \mathcal{U}(x), \quad t \geq 0. \quad (2)$$

The next subsection describes the linearized dynamics. Sections 3.2–3.3 provide an overview of design techniques for the parametrized transition family $\{P_\zeta : \zeta \in \mathbb{R}\}$, to ensure that the linearized input-output model has desirable properties for control at the grid level.

Linearized mean-field model The mean-field model (1) is a state space model that is linear in the state μ_t , and nonlinear in the input ζ_t . The observation equation (2) is also linear as a function of the state. Assumptions imposed in [11, 13, 36] imply that the input is a continuous function of these values. The design of the feedback law $\zeta_t = \phi_t(y_0, \dots, y_t)$ is based on a linearization of this state space model.

The linearized input-output model requires additional notation. The derivative of the transition matrix is also a $d \times d$ matrix, denoted

$$\mathcal{E}_\zeta = \frac{d}{d\zeta} P_\zeta \quad (3)$$

Denote $\widetilde{\mathcal{U}}_\zeta = \mathcal{U} - \bar{\mathcal{U}}_\zeta$, with $\bar{\mathcal{U}}_\zeta = \pi_\zeta(\mathcal{U})$.

The invariant pmf π_ζ for P_ζ is regarded as the equilibrium state for the mean-field model (1), with respect to the constant input value $\zeta_i \equiv \zeta$. The linearization about this equilibrium is described in Prop. 1. The proof can be found in [36, Prop. 2.4].

Proposition 1. *Consider a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$ that are continuously differentiable in ζ . Assume also that P_ζ is irreducible and aperiodic for each ζ . The unique invariant pmf π_ζ is an equilibrium for (1) when ζ takes on this constant value. The input-output model with state evolution (1), input ζ , and output (2) admits a linearization about this equilibrium. It is described as a d -dimensional state space model with transfer function,*

$$G_\zeta(z) = C[Iz - A]^{-1}B \quad (4)$$

in which $A = P_\zeta^i$, $C_i = \widetilde{\mathcal{U}}_\zeta(x^i)$ for each i , and

$$B_i = \sum_x \pi_\zeta(x) \mathcal{E}_\zeta(x, x^i), \quad 1 \leq i \leq d \quad (5)$$

3.2 Local control design

It is assumed throughout this chapter that the family of transition matrices used for distributed control is of the form,

$$P_\zeta(x, x') := P_0(x, x') \exp(h_\zeta(x, x') - \Lambda_{h_\zeta}(x)) \quad (6)$$

in which h_ζ is continuously differentiable in ζ , and Λ_{h_ζ} is the normalizing constant

$$\Lambda_{h_\zeta}(x) := \log\left(\sum_{x'} P_0(x, x') \exp(h_\zeta(x, x'))\right) \quad (7)$$

Each P_ζ is irreducible and aperiodic under the assumption that this is true for P_0 .

Myopic design and the exponential family. A simple choice is the *myopic design*. This is obtained by setting $h_\zeta(x, x') = \zeta \mathcal{U}(x')$,

$$P_\zeta^{myop}(x, x') := P_0(x, x') \exp(\zeta \mathcal{U}(x') - \Lambda_\zeta(x)) \quad (8)$$

with the normalizing constant $\Lambda_\zeta(x) := \log\left(\sum_{x'} P_0(x, x') \exp(\zeta \mathcal{W}(x'))\right)$. This corresponds to a tilted probability transition matrix, favoring the transitions to states with lower power consumption when $\zeta < 0$, and to states with higher power consumption when $\zeta > 0$.

Advantages of this design include ease of implementation, and the straightforward generalization to the continuous state space case. This generalization will be illustrated in Sections 4 and 6.

It is possible to consider any other family of functions, linear with respect to ζ , leading to an exponential family for $\{P_\zeta : \zeta \in \mathbb{R}\}$,

$$h_\zeta(x, x') = \zeta H_0(x, x'). \quad (9)$$

The choice of H_0 will typically correspond to the linearization of a more advanced design around the value $\zeta = 0$ (or some other fixed value of ζ). One example is given in Section 3.3.

Individual Perspective Design. Consider a finite-time-horizon optimization problem: For a given terminal time T , let p_0 denote the pmf on strings of length T :

$$p_0(x_1, \dots, x_T) = \prod_{i=0}^{T-1} P_0(x_i, x_{i+1}),$$

where $x_0 \in \mathsf{X}$ is assumed to be given. The scalar $\zeta \in \mathbb{R}$ is interpreted as a weighting parameter in the following definition of total welfare. For any pmf p , this is defined as the weighted difference,

$$\mathcal{W}_T(p) = \zeta \mathbb{E}_p \left[\sum_{t=1}^T \mathcal{W}(X_t) \right] - D(p \| p_0) \quad (10)$$

where the expectation is with respect to p , and D denotes relative entropy:

$$D(p \| p_0) := \sum_{x_1, \dots, x_T} \log \left(\frac{p(x_1, \dots, x_T)}{p_0(x_1, \dots, x_T)} \right) p(x_1, \dots, x_T)$$

It is easy to check that the myopic design is an optimizer for the horizon $T = 1$,

$$P_\zeta^{\text{myop}}(x_0, \cdot) \in \arg \max_p \mathcal{W}_1(p).$$

The infinite-horizon mean welfare is denoted,

$$\eta_\zeta^* = \lim_{T \rightarrow \infty} \frac{1}{T} \mathcal{W}_T(p_T^*) \quad (11)$$

The two terms in the welfare function (10) represent the two conflicting goals: To provide service to the grid and to reduce deviation of the load's behavior from the

nominal. If the controlled probability p is chosen to be different from p_0 , it potentially reduces the QoS to the consumer, which is modeled by the term “ $-D(p||p_0)$.”

Recall that $\mathcal{U}(X_t)$ is equal to the power consumption of the load at time t . If the grid operator desires lower power demand than the nominal value, this goal is modeled through the first term in (10) whenever the parameter ζ is negative.

A solution to the infinite horizon problem is given by a time-homogenous Markov chain whose transition matrix is obtained following the solution of an eigenvector problem, based on the $d \times d$ matrix,

$$\widehat{P}(x, x') = \exp(\zeta \mathcal{U}(x)) P_0(x, x'), \quad x, x' \in \mathbf{X}. \quad (12)$$

Let $\lambda > 0$ denote the Perron-Frobenius eigenvalue, and v the eigenvector with non-negative entries satisfying,

$$\widehat{P}v = \lambda v \quad (13)$$

The proof of Prop. 2 is contained in [36, Section II], following [44].

Proposition 2. *If P_0 is irreducible, an optimizing p^* that achieves (11) is defined by a time-homogeneous Markov chain whose transition probability is defined by,*

$$\check{P}_\zeta(x, x') = \frac{1}{\lambda} \frac{1}{v(x)} \widehat{P}(x, x') v(x'), \quad x, x' \in \mathbf{X}. \quad (14)$$

3.3 Uncontrolled dynamics

In many cases it is not possible to apply the IPD solution in the form (14) because a portion of the stochastic dynamics are not directly controllable. Consider a load model in which the full state space is the Cartesian product $\mathbf{X} = \mathbf{X}^u \times \mathbf{X}^n$, where \mathbf{X}^u are components of the state that can be directly manipulated through control.

In prior work [7, 6], the following conditional-independence structure is assumed: for each state $x = (x_u, x_n)$, and each $\zeta \in \mathbb{R}$,

$$\begin{aligned} \check{P}_\zeta(x, x') &= R_\zeta(x, x'_u) Q_0(x, x'_n), \\ R_\zeta(x, x'_u) &= R_0(x, x'_u) \exp(h_\zeta(x, x'_u) - \Lambda_{h_\zeta}(x)) \end{aligned} \quad (15)$$

where $\sum_{x'_u} R_\zeta(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) = 1$ for each x and ζ . The matrix Q_0 is out of our control – this models load dynamics and exogenous disturbances. For example, it may be used to model the impact of the weather on the climate of a building. The matrices $\{R_\zeta\}$ are a product of design.

It is reasonable to assume that \mathcal{U} is a function only of \mathbf{X}^u ; the power state is directly controllable. In this case the myopic design (8) is unchanged, $h_\zeta(x, x'_u) = \zeta \mathcal{U}(x'_u)$.

The formulation of the IPD optimization problem is unchanged, but its solution is not in the form (14). A computational ODE approach is introduced in [7, 6]: for a

vector field \mathcal{V} whose domain and range are functions on $\mathbb{X} \times \mathbb{X}^u$,

$$\frac{d}{d\zeta} h_\zeta = \mathcal{V}(h_\zeta), \quad \zeta \in \mathbb{R}, \quad h_0 \equiv 1.$$

Besides its computational value, this approach provides a useful alternative to the myopic design. The function $H_0 = \mathcal{V}(h_0)$ can be used in the exponential family design (9). It is shown in [7] that this function is a solution to Poisson's equation for the nominal model: $P_0 H_0 = H_0 - \tilde{\mathcal{W}}_0$.

Motivation for the IPD design or its exponential family approximation is in part empirical. In nearly every numerical experiment conducted to-date, it is found that the resulting input-output mean field model appears nearly linear over a large range of ζ , and also minimum phase. Moreover, in nearly all cases the linearization (4) is *passive* when the delay is removed. That is, the transfer function $zC[Iz - A]^{-1}B$ is strictly positive real.

Passivity can be established mathematically for a restricted class of models [5], or using a different ODE called the system perspective design (SPD) [6].

3.4 Quality of service and opt-out

In analysis of QoS it is convenient to consider a steady-state setting: the state process for each load is assumed to be a stationary process on the two-sided time interval. It is also useful to consider a functional form for QoS – the following conventions were introduced in [11].

Several QoS metrics may be considered simultaneously, but each are assumed to be of the following form. Assumed given is a function $\ell: \mathbb{X} \rightarrow \mathbb{R}$, defined so that $L_t^i := \ell(X_t^i)$ describes a “snap-shot” indication of QoS for the i th load at time t . The function ℓ may represent the temperature of a TCL, cycling of an on/off load, or power consumption as a function of $x \in \mathbb{X}$.

Second is a stable transfer function denoted H_ℓ . The QoS of the i th load at time t is defined by passing L^i through the transfer function H_ℓ . Two classes of transfer functions H_ℓ are considered in prior research and examples in this chapter:

- (i) Summation over a finite time horizon T_f :

$$\mathcal{L}_t^i = \sum_{k=0}^{T_f} \ell(X_{t-k}^i). \quad (16)$$

- (ii) Discounted sum, with discount factor $\beta \in [0, 1)$:

$$\mathcal{L}_t^i = \sum_{k=0}^{\infty} \beta^k \ell(X_{t-k}^i). \quad (17)$$

When β is close to unity, or T_f is very large, then these QoS metrics can be approximated by Gaussian random variable by appealing to the Central Limit Theorem [11]. A Gaussian distribution indicates that QoS for some individuals in the population will sometimes take on unacceptable values.

QoS can be constrained by imposing an additional layer of control at each load. A simple mechanism is *opt-out control*.

The opt-out mechanism is based on pre-defined upper and lower limits, denoted b_+ and b_- . A load ignores a command from grid operator if it will result in $\mathcal{L}_{t+1}^i \notin [b_-, b_+]$, and takes an alternative action so that $\mathcal{L}_{t+1}^i \in [b_-, b_+]$. This ensures that the QoS metric of each load remains within the predefined interval for all time.

Numerical examples are presented in [11] for both residential pools and TCLs. Some of these results are surveyed in Section 5. Negative impact on tracking performance is observed in numerical experiments only when the QoS interval $[b_-, b_+]$ is small (e.g., b_+ is less than the mean plus one standard deviation of the distribution without opt-out control).

4 Example: Thermostatically Controlled Loads

This special case is dominant in much of the literature on demand dispatch. Examples of thermostatically controlled loads (TCLs) include refrigerators, water heaters and air-conditioning. Each of these loads is already equipped with primitive “local intelligence” based on a *deadband* (or *hysteresis interval*): there is a sensor that measures the temperature of the unit, and turns the power on when the measured value reaches one end of this deadband.

The state process for a TCL at time t will be of the form

$$X(t) = (X_u(t), X_n(t)) = (m(t), \Theta(t)), \quad (18)$$

in which $m(t) \in \{0, 1\}$ denotes the power mode (the value “1” indicating the unit is on), and $\Theta(t)$ the inside temperature of the load. Exogenous disturbances that directly influence Θ include ambient temperature, and usage: the inside temperature of a refrigerator is impacted by an open door, and the temperature of water in a water heater is influenced by the rate of flow of water out of the unit.

The remainder of this section is restricted to a residential water heater (WH). This will simplify discussion, and extensions to other TCLs are often straightforward.

Nominal model. The standard ODE model of a water heater is the first-order linear system:

$$\frac{d}{dt}\Theta(t) = -\lambda[\Theta(t) - \Theta^a(t)] + \gamma m(t) - \alpha[\Theta(t) - \Theta^{in}(t)]f(t), \quad (19)$$

for constants $(\lambda, \gamma, \alpha)$, in which $\Theta(t)$ is the temperature of the water in the tank, $\Theta^a(t)$ is ambient temperature, $\Theta^{in}(t)$ is temperature of the cold water entering the tank (degrees Fahrenheit), $f(t)$ is flow rate of hot water from the WH (gallons/s), and

$m(t)$ is the power mode of the WH (“on” indicated by $m(t) = 1$). The corresponding power consumed by a WH when $m(t) = 1$ is denoted P_{on} .

The upper and lower temperature limits that define the deadband are denoted Θ_- , Θ_+ , respectively. A standard residential water heater in the U.S. has the following typical behavior: At the moment that $\Theta(t)$ reaches the lower limit Θ_- , the unit turns on, and remains on until the time t_+ at which $\Theta(t_+) = \Theta_+$. The unit then turns off and begins to cool. It may take 6 hours to once again reach the lower limit, while the time to heat the water is much shorter.

The nominal model used for local control design is based on an approximation of this typical behavior, in which with some probability the unit turns on before $\Theta(t)$ reaches Θ_- , and the unit may also turn off before reaching the maximum temperature Θ_+ . The definition of the nominal model is based on the specification of two cumulative distribution functions (CDFs) for the temperature at which the load turns on or turns off, denoted F^\oplus and F^\ominus . Random variables with these CDFs are denoted $\tilde{\Theta}^\oplus$ and $\tilde{\Theta}^\ominus$, so that

$$F^\oplus(\theta) = \text{P}\{\tilde{\Theta}^\oplus \leq \theta\}, \quad F^\ominus(\theta) = \text{P}\{\tilde{\Theta}^\ominus \leq \theta\}, \quad \theta \in \mathbb{R}.$$

It is always assumed that $\tilde{\Theta}^\oplus$ and $\tilde{\Theta}^\ominus$ take values in the interval $[\Theta_-, \Theta_+]$, which implies that $F^\oplus(\theta) = F^\ominus(\theta) = 1$ for $\theta \geq \Theta_+$ and $F^\oplus(\theta) = F^\ominus(\theta) = 0$ for $\theta < \Theta_-$.

A particular design for F^\ominus is obtained on fixing three parameters $\theta_0^\ominus \in [\Theta_-, \Theta_+]$, and constants $\rho \in (0, 1)$ and $\kappa > 1$:

$$F^\ominus(\theta) = (1 - \rho) \frac{[\theta - \theta_0^\ominus]_+^\kappa}{[\Theta_+ - \theta_0^\ominus]^\kappa}, \quad \theta \in [\Theta_-, \Theta_+],$$

where $[x]_+ := \max(0, x)$ for $x \in \mathbb{R}$. In a symmetric model, the other CDF is defined by the transformation,

$$F^\oplus(\theta) = 1 - \lim_{\theta' \downarrow \theta} F^\ominus(\Theta_+ + \Theta_- - \theta')$$

Fig. 4 illustrates a particular special case of the symmetric model.

It is assumed that the local control operates in discrete-time. By choice of time-units, without loss of generality it is assumed that the sampling interval is 1 unit. At time instance k , if the water heater is on (i.e., $m(k) = 1$), then it turns off at time $k + 1$ with probability,

$$p^\ominus(k+1) = \frac{[F^\ominus(\Theta(k+1)) - F^\ominus(\Theta(k))]_+}{1 - F^\ominus(\Theta(k))}$$

If $\Theta(k+1) \leq \Theta(k)$, then this probability is zero. Similarly, if the load is off, then it turns on with probability


$$p^\oplus(k+1) = \frac{[F^\oplus(\Theta(k)) - F^\oplus(\Theta(k+1))]_+}{F^\oplus(\Theta(k))}$$

The nominal behavior of the power mode can be expressed

$$\begin{aligned} \mathbb{P}\{m(k) = 1 \mid \theta(k-1), \theta(k), m(k-1) = 0\} &= p^\oplus(k) \\ \mathbb{P}\{m(k) = 0 \mid \theta(k-1), \theta(k), m(k-1) = 1\} &= p^\ominus(k) \end{aligned} \quad (20)$$

The IPD and SPD designs were obtained in [6] based on a similar nominal model for a residential refrigerator.

Fig. 4 Nominal model for a water heater: an instance of the symmetric model.



The myopic design (15) is obtained through an exponential tilting:

$$p_\zeta^\oplus(k) := \frac{p^\oplus(k)e^\zeta}{p^\oplus(k)e^\zeta + 1 - p^\oplus(k)}, \quad p_\zeta^\ominus(k) := \frac{p^\ominus(k)}{p^\ominus(k) + (1 - p^\ominus(k))e^\zeta}$$

If $p^\oplus(k) > 0$, then the probability $p_\zeta^\oplus(k)$ is strictly increasing in ζ , approaching 1 as $\zeta \rightarrow \infty$; it approaches 0 as $\zeta \rightarrow -\infty$, provided $p^\oplus(k) < 1$.

System identification. Power, temperature, and usage data from residential water heaters was obtained through our partners at ORNL.⁴ The constants $(\lambda, \gamma, \alpha)$ were estimated using least squares. The parameter values listed in Table 1 reflect the range of values observed in actual data.

A testbed was created to simulate a collection of $N = 100,000$ water heaters with usage. Each evolves according to the ODE (19), but parameters were different for each of the N loads: parameters were chosen via uniform sampling of the values in Table 1. A simulation model for usage at each load was created, based on sampling from historical usage of actual water heaters.


The mean-field model is a nonlinear input-output system with input ζ and output equal to power deviation, \mathbf{y} . An approximate linear model was obtained through least squares, in which the input ζ was taken to be the swept-sine: $\zeta(t) = 1.5 \sin(10^{-7}t^2)$ for $0 \leq t \leq 432 \times 10^5$ sec. (5 days). Fig. 5 shows results from the estimation experiment for two different model orders. The Bode plots shown represent the approximate model in continuous time. The 5th order model predicts that the gain of the linearization vanishes as the frequency tends to zero (DC). This is a physical reality for this example.

⁴ Water heater data provided by Ecotope, Inc., with funding from the Northwest Energy Efficiency Alliance (NEEA) and the Bonneville Power Administration (BPA).

Temp. Ranges	ODE Pars.	Loc. Control
$\Theta_+ \in [118, 122]$ F	$\lambda \in [8, 12.5] \times 10^{-6}$	$T_s = 15$ sec
$\Theta_- \in [108, 112]$ F	$\gamma \in [2.6, 2.8] \times 10^{-2}$	$\kappa = 4$
$\Theta^a \in [68, 72]$ F	$\alpha \in [6.5, 6.7] \times 10^{-2}$	$\rho = 0.8$
$\Theta^{in} \in [68, 72]$ F	$P_{on} = 4.5$ kW	$\theta_0 = \Theta_-$

Table 1 Parameters for nominal model for water heaters.

The linearization is minimum phase and stable. Its gain is approximately constant in the frequency range $[5 \times 10^{-4}, 10^{-2}]$ rad/s. It is expected that a collection of water heaters can accurately track signals in this frequency range.

**Fig. 5** Least square estimates of the transfer function for water heaters.

Tracking. Design at the macro level is most easily performed for a model in continuous time. A PI controller $G_c(s) = K_P + K_I/s$ was designed based on the linearized mean-field model. The values $K_P = 10^5$ and $K_I = 500$ result in a crossover frequency $\omega_c = 0.03$ rad/s (corresponding to a time period of approximately 3.5 minutes), with a 75° phase margin.

The balancing reserves signal from the Bonneville Power Administration (BPA) was used in the tracking experiments described in this section. A typical windy day, February 19, 2016, was chosen for the experiments described here. The signal was filtered using a second-order Butterworth high pass filter with a cut-off frequency at 8×10^{-4} rad/s (corresponding to a sine wave with period of approximately 2 hours).

Fig. 6 shows results from several numerical experiments. The three rows are differentiated by the regulation signal: In the first row $r \equiv 0$, in the second the absolute value of the regulation signal takes a maximum value of about 8 MW, and in the final row the prior regulation signal was multiplied by 4. Exact tracking is not feasible over the entire period for the largest regulation signal (results shown in the bottom left plot), but the performance remains nearly perfect over time periods for which $|r_t|$ does not exceed about 90% of the nominal power consumption.

The second column shows evolution of temperature and the power mode for a typical load in the three cases. The seed for the random number generator was iden-

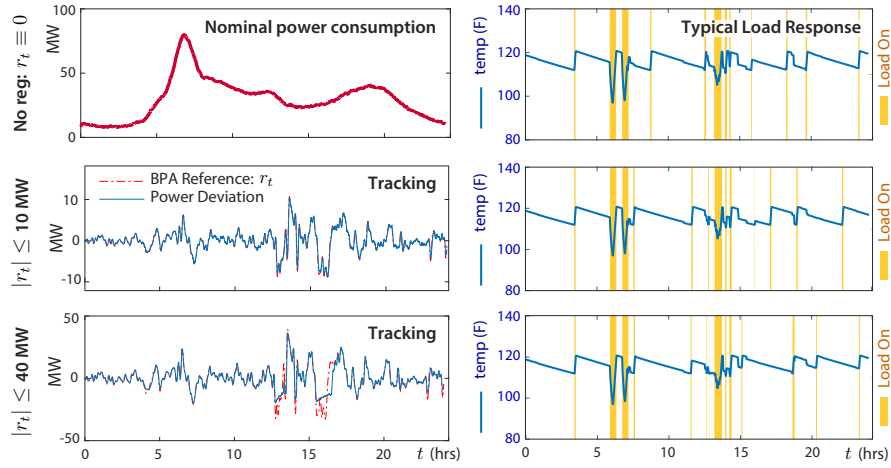


Fig. 6 Tracking results with 100,000 water heaters, and the behavior of a single water heater in three cases, distinguished by the reference signal r . The morning peak in nominal power consumption is consistent with typical water usage included in the simulation experiments.

tical in each of the three experiments. It is amazing to see that the evolution of temperature and power mode is hardly impacted by local control.

These loads are equally valuable for contingency and ramping services. Fig. 7 shows recent results that illustrate the potential. In these experiments the water flow was set to zero; in this case, the nominal power consumption for 100,000 loads is approximately 8 MW. Each plot is a particular saw-tooth wave, scaled to reach the maximum lower limit of -8 MW.

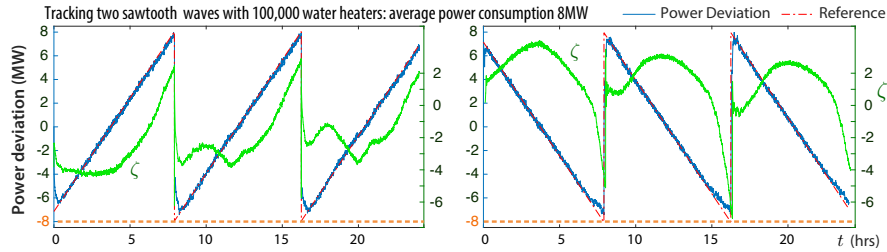


Fig. 7 Tracking a pair of saw tooth waves with 10^5 water heaters.

5 Example: Residential Pools

The paper [36] and its sequels [11, 13, 12] consider this system architecture in which the loads are a collection of pools. The motivation for considering pools is the in-

herent flexibility of pool cleaning, and because the total load in a region can be very large. The maximum load is approximately 1 GW in California or Florida.

The state space for the discrete-time model is the finite set: $X = \{(m, j) : m \in \{0, 1\}, j \in \{1, 2, \dots, \mathcal{J}\}\}$. The first variable indicates the power mode, with $m = 1$ indicating the power is on. The second integer component is interpreted as follows: The load samples the grid signal periodically (the sampling increments are assumed to be deterministic, or i.i.d. and distributed according to a geometric distribution). At the time of the t -th sample, if $X_t^i = (0, j)$ then the load has remained off for the past j sampling times, and was turned off at sampling time $t - j$; the interpretation of $X_t^i = (1, j)$ is symmetrical.

A nominal model can be constructed in a manner similar to the case of TCLs. In this application, each CDF models the time at which the power mode changes. The IPD solution obtained using Prop. 2 is considered in [36] for a model without geometric sampling, and in [13] for the present model. The linearized mean-field models obtained in [36, 13] are minimum phase and have a resonance at a frequency corresponding to a period of approximately 24 hours.

The numerical results that follow are based on a stochastic simulation of a large number of pools. Each pool consumes 1 kW when in operation. Both 12 and 8 hour nominal daily cleaning cycles are considered. Tracking results with a heterogeneous population of loads are described in [36, 13].




Fig. 8 Tracking is nearly perfect with reference scaled to the theoretical limit.

Tracking and contingency reserves. The first set of experiments concern tracking of the balancing reserves deployed over one week at BPA. The sampling time is taken to be every five minutes. The signal was filtered to remove the highest frequency components. Tracking the original signal is possible, but with reduced overall capacity [10]. In each example, the number of loads is equal to $N = 10^4$.

A theoretical limit on capacity is obtained by considering the fraction of pool pumps that are operating in nominal steady-state:

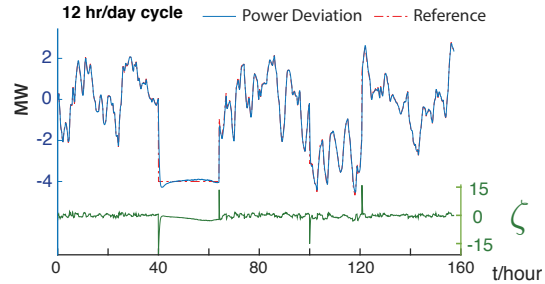
$$\pi_0^\oplus = \sum_i \pi_0(1, i) \quad \text{where } \pi_0 \text{ is invariant for } P_0.$$

Upper and lower bounds on power deviation are defined as follows, in units of kW:

$$\{+Demand^*, -Supply^*\} := \{(1 - \pi_0^\oplus) \times N, -\pi_0^\oplus \times N\}$$

This is approximately $\{+5, -5\}$ MW for 12 hr/day cycling, and $\{+6.6, -3.4\}$ MW for 8 hr/day. Results from simulation experiments shown in Fig. 8 show that these limits are nearly attainable in each case.

Fig. 9 Range of services provided by pools: contingency reserves and balancing can be supplied simultaneously.



The potential for virtual energy storage goes far beyond tracking a balancing reserve signal. Experiments were conducted in [10] to investigate the potential for providing contingency reserves in conjunction with balancing reserves. A reference signal was constructed based on the one used in the previous experiments, with two changes: during the period $[40, 64]$ hours, the reference signal was *replaced* by a 4 MW power reduction and during the period $[100, 124]$ hours the reference signal was *augmented* with a 3 MW power reduction. The PI control parameters were unchanged. Fig. 9 shows again nearly perfect tracking.

In practice, the signal ζ should be transformed so that it is zero energy over the week — this will help to ensure that QoS constraints are not violated.




Fig. 10 Improvement of QoS with the introduction of local opt-out control.

Quality of service and opt-out. The grid is receiving nearly perfect services – what about the service offered by each load to its owner?

In the experiments conducted to produce either of the plots in Fig. 8, a histogram of total operation hours over the time horizon appears approximately Gaussian with mean value 78 hours (consistent with the 12hrs/day cleaning cycle for each pool.) The Gaussian approximation can be used to estimate the fraction of pools that are over-cleaned or under cleaned over the week [12].

To investigate the impact of opt-out control on QoS and capacity, consider a family of models parameterized by $0 \leq \varepsilon \leq 1$. The reference signal is obtained by scaling: $r_t^\varepsilon = \varepsilon r_t^1$, where r^1 is the reference signal considered in prior experiments.

Consider the following two QoS metrics suitable for this application:

- (i) Normalized power consumption of a load, $\ell(X_t^i) = \mathcal{W}(X_t^i) - \bar{y}^0$, where \bar{y}^0 is the nominal steady-state mean (under $\zeta = 0$).
- (ii) On/off cycling:

$$\ell^c(X_t^i, X_{t+1}^i) = \left| \sum_j \left(\mathbf{1}\{X_{t+1}^i = (1, j)\} - \mathbf{1}\{X_t^i = (1, j)\} \right) \right| \quad (21)$$

The discounted sum (17) was used to define \mathcal{L}_t^i in the experiments surveyed here. The discount factor $\beta = 2779/2880$ was chosen so that the discounted sum (17) is similar to the moving window QoS metric (16) with T_f corresponding to 10 days (recall the sampling period is five minutes).

Fig. 10 illustrates an example of QoS improvement based on a 15% constraint on both QoS metrics, using the reference signal r^1 . The opt out rate is very small in this case (much smaller than predicted by the corresponding tails of the histogram without opt-out), and the tracking is nearly perfect.


Four QoS intervals were considered corresponding to constraints of, respectively, 5%, 10%, 15%, and 20%. For example, a cleaning QoS constraint of 5% corresponds to ± 3 cleaning hours — a very tight constraint over a 10-day time horizon. No lower bound was imposed on cycling QoS.

A normalized root mean square error (NRMSE) was adopted as the metric for grid level tracking performance:

$$\text{NRMSE} = \frac{1}{\varepsilon} \frac{\text{RMS}(e) - \text{RMS}(e^0)}{\text{RMS}(r^1)}, \quad (22)$$

where e, e^0 are error signals with and without the reference input, and $\text{RMS}(f) = \sqrt{\frac{1}{T} \sum_{k=1}^T f_k^2}$ for any signal f over a time horizon T .

Fig. 11 Tracking performance with two QoS constraints — total cleaning hours and cycling. Opt-out control has little impact on tracking error over a large range of opt-out intervals.



Tracking performance for a range of opt-out parameters is summarized in Fig. 11 using 16 colored bars, distinguished by each pair of QoS constraints. Each bar represents tracking errors for different reference signal scaling factors, $0.1 \leq \varepsilon \leq 1$. The darkest color represents NRMSE (22) of 10% or greater, and lighter colors represent smaller values of NRMSE (indicated on the color bar label). This shows that opt-out control based on the two QoS metrics has little impact on tracking error over a large range of opt-out intervals.

6 Example: Battery Systems


Many believe that there will be a battery revolution over the next decade – small battery systems will be distributed across the grid at residential homes, and in racks at gas stations where owners of electric vehicles can exchange their old battery.

Coordination of the population can be performed as described for residential loads, even though the physics and QoS constraints are very different. Longevity of a battery requires constraints on the state of charge (SoC), as well as ramping and temperature constraints.

A demand dispatch architecture is proposed in [4], in which the state space for an individual battery is again a Cartesian product: a particular state is denoted $x = (m, s)$, where $m \in \{\text{ch}, \text{dis}, \text{id}\}$ denotes charging mode, and $s \in [0, 1]$ denotes the SoC. The power delivery at state x depends only on charging mode: $\mathcal{U}(\text{ch}, s) = \mathcal{U}_{\text{ch}} < 0$, $\mathcal{U}(\text{id}, s) = 0$, $\mathcal{U}(\text{dis}, s) = \mathcal{U}_{\text{dis}} > 0$.

The design of the family of transition matrices $\{P_\zeta\}$ on the state space $X = \{\text{ch}, \text{dis}, \text{id}\} \times [0, 1]$ is based on the myopic policy. The main difficulty compared to loads is that there is no obvious nominal model P_0 (for loads, this is taken as a stochastic perturbation of a deterministic model for $\zeta \equiv 0$.) The nominal model P_0 for batteries was chosen so that the invariant pmf π_0 would have most of its mass concentrated at SoC near 60%. The randomized decision rule is designed to encourage idle time for each battery, and to avoid extreme SoC levels and frequent switching of modes.

Fig. 12 Design of switching probability functions for the battery system. A weighted coin is flipped to determine if the battery will stay in its current power mode.



Let $X_t^i = (M_t^i, S_t^i)$ denote the state of i th battery at time t . The SoC evolves as a controlled random walk: $S_{t+1}^i = S_t^i + h\delta_{\text{ch}}$, if $M_t^i = \text{ch}$, $S_{t+1}^i = S_t^i - h\delta_{\text{dis}}$, if $M_t^i = \text{dis}$, $S_{t+1}^i = S_t^i$, if $M_t^i = \text{id}$, where h is the time step length, and δ_{ch} and δ_{dis} charging and discharging rates. The dynamics of the first component are governed by a “two coin-flip” randomized policy: In state (m, s) , a weighted coin is flipped to determine if the battery will stay in its current power mode. The design of the probability functions $p_{\text{ch}}, p_{\text{dis}}, p_{\text{id}} : [0, 1] \rightarrow [0, 1]$ that model the probability to stay in the charging, discharging or idle mode respectively are shown in Fig. 12. If the outcome of the first coin flip is “mode change”, then a second coin flip is used to decide which of the remaining two modes the battery is going to switch to. This choice is done with the probabilities proportional to the values of the p -functions of the alternative power modes. For example, in state (ch, s) , the battery changes its mode to idle with probability $(1 - p_{\text{ch}}(s)) \times p_{\text{id}}(s) / (p_{\text{id}}(s) + p_{\text{dis}}(s))$

The nominal design in Fig. 12 was chosen by setting a *target SoC interval* to 40 - 80% SoC (to allow ramping capability while avoiding extreme SoC levels): If the battery is charging, it will remain charging with probability almost 1 until it reaches 40% SoC. The probability to keep charging then decreases and reaches almost 0 at 80% SoC. The design of p_{dis} is symmetrical. The function p_{id} has values almost 1 for 50-70% SoC values and it is almost 0 outside the target interval.

Batteries are ideal for tracking signals of higher frequency – timescales of tens of seconds to many minutes. An example is the RegD signal used at PJM. It is found that tracking of this signal is nearly perfect using a combination of local control at the battery, and a PI compensator at the BA (see [4] for details).

It might be assumed that the randomized control law would lead to excess cycling of batteries. In fact, the behavior of a typical battery behaved nearly deterministically. Typical behavior is illustrated in Fig. 13.

A measure of battery activity is *mileage* [37], which is regarded here as an example of a QoS metric. For a time horizon T , mileage for battery k is denoted

$$\mathcal{L}_T^k = \sum_{t=1}^T |\mathcal{W}(X_t^k)|, \quad (23)$$

and \mathcal{L}_T denotes the average over $k = 1, \dots, N$. The *excess operation* is the normalized difference:

$$\mathcal{O}_T = \frac{\mathcal{L}_T - \mathcal{L}_T^*}{\mathcal{L}_T^*}, \quad \text{where } \mathcal{L}_T^* = \frac{1}{N} \sum_{t=1}^T |r_t|$$

It can be shown using Jensen's inequality that $\mathcal{O}_T \geq 0$ in the ideal case of perfect tracking. If *each* battery tracks the reference exactly, then $\mathcal{O}_T = 0$. In numerical experiments the value $\mathcal{O}_T \approx 3\%$ is typical.




Fig. 13 Left: Tracking the PJM RegD test signal with $N = 10^3$ batteries. Right: comparison of SoC of an individual and the average of the population.

7 Conclusions

With appropriate filtering and local control, loads can provide excellent grid services without two-way communication. While there is some cost to install hardware on appliances that can receive a signal from a balancing authority, in the long run this will be far less costly than batteries.

The numerical results presented in this paper, in particular the tracking results illustrated in Figures 6, 7 and 9, show that VES working in conjunction with traditional resources can provide balancing services, ramping services and contingency reserves simultaneously. It is likely that water heaters, pools and agriculture loads in California can provide the resources required to address their future grid service requirements.

Current research questions include:

- (i) The application of reinforcement learning may be valuable for learning the local control law, such as an extension of Z-learning [44] to the IPD approach.
- (ii) The numerical results presented here concern signals on time scales of tens of seconds and slower. Ancillary service on faster time-scales correspond to what is called primary reserves. Control design requires more care in this context because poor performance can induce grid instability [30, 31].
- (iii) Further research is required to better estimate capacity in terms of both energy and power [21].
- (iv) The impact of usage is not entirely understood. Numerical results presented in Section 4 suggest that this is not an obstacle in the case of water heaters. Air-conditioning is a greater challenge because variations in load are much greater.
- (v) A question posed in [31]: Does the load need to receive a signal from the BA? It is possible that some VES resources can provide valuable services using only local measurements. Frequency (as well as voltage) measurements can be obtained inexpensively at loads, and these measurements are similar to those used by the BA to construct analogs of our “ ζ ” today. The advantage of distributed control is reduced cost due to reduced communication between a BA and loads.

The BA will continue to regulate tie-line error, and they will continue to regulate frequency as well. It is hoped that the balancing resources required at the BA will be reduced through this extra layer of distributed control.

References

1. 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.
2. D. S. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. *Math. Oper. Res.*, 27(4):819–840, 2002.

3. J. Y. L. Boudec, D. McDonald, and J. Munding. A generic mean field convergence result for systems of interacting objects. In *Fourth International Conference on the Quantitative Evaluation of Systems (QEST 2007)*, pages 3–18, Sept 2007.
4. A. Bušić, M. U. Hashmi, and S. Meyn. Distributed control of a fleet of batteries. In *American Control Conference*, pages 3406–3411, May 2017.
5. A. Bušić and S. Meyn. Passive dynamics in mean field control. In *Proc. 53rd IEEE Conference on Decision and Control*, pages 2716–2721, Dec 2014.
6. A. Bušić and S. Meyn. Distributed randomized control for demand dispatch. In *IEEE Conference on Decision and Control*, pages 6964–6971, Dec 2016.
7. A. Bušić and S. Meyn. Ordinary Differential Equation Methods For Markov Decision Processes and Application to Kullback-Leibler Control Cost. *Under revision, SIAM J. Control and Opt.*, 2016.
8. California ISO – Folsom, CA 95763-9014. ISO Today. Online www.ca.iso.com/Pages/TodaysOutlook.aspx.
9. D. Callaway and I. Hiskens. Achieving controllability of electric loads. *Proceedings of the IEEE*, 99(1):184–199, January 2011.
10. Y. Chen. *Markovian demand dispatch design for virtual energy storage to support renewable energy integration*. PhD thesis, University of Florida, Gainesville, FL, USA, 2016.
11. Y. Chen, A. Bušić, and S. Meyn. Estimation and control of quality of service in demand dispatch. *IEEE Trans. on Smart Grid*, PP(99):1–1, 2017.
12. Y. Chen, A. Bušić, and S. Meyn. Ergodic theory for controlled Markov chains with stationary inputs. *ArXiv e-prints and to appear in Annals of Applied Prob.*, April 2017.
13. 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.
14. M. Chertkov and V. Y. Chernyak. Ensemble control of cycling energy loads: Markov Decision Approach. In *IMA volume on the control of energy markets and grids*. Springer, 2017.
15. 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.
16. K. Christakou, D.-C. Tomozei, J.-Y. Le Boudec, and M. Paolone. GECN: primary voltage control for active distribution networks via real-time demand-response. *IEEE Trans. on Smart Grid*, 5(2):622–631, March 2014.
17. P. Fairley. Energy storage: Power revolution. *Nature*, 526:S102–S104, October 29 2015.
18. W. H. Fleming and S. K. Mitter. Optimal control and nonlinear filtering for nondegenerate diffusion processes. *Stochastics*, 8(1):63–77, 1982.
19. H. Hao, Y. Lin, A. Kowli, P. Barooah, and S. Meyn. Ancillary service to the grid through control of fans in commercial building HVAC systems. *IEEE Trans. on Smart Grid*, 5(4):2066–2074, July 2014.
20. H. Hao, T. Middelkoop, P. Barooah, and S. Meyn. How demand response from commercial buildings will provide the regulation needs of the grid. In *50th Allerton Conference on Communication, Control, and Computing*, pages 1908–1913, 2012.
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. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.*, 6(3):221–251, 2006.
23. L. P. Kadanoff. More is the same; phase transitions and mean field theories. *J. Stat. Phys.*, 137(5-6):777–797, 2009.
24. M. Kárný. Towards fully probabilistic control design. *Automatica*, 32(12):1719–1722, 1996.
25. A. Kizilkale and R. Malhame. Mean field based control of power system dispersed energy storage devices for peak load relief. In *52nd IEEE Conference on Decision and Control (CDC)*, pages 4971–4976, 2013.
26. A. Kizilkale and R. Malhame. A class of collective target tracking problems in energy systems: Cooperative versus non-cooperative mean field control solutions. In *IEEE Conference on Decision and Control*, pages 3493–3498, 2014.

27. J.-M. Lasry and P.-L. Lions. Mean field games. *Jpn. J. Math.*, 2(1):229–260, 2007.
28. O. Madani, S. Hanks, and A. Condon. On the undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence*, 147(1-2):5–34, 2003.
29. R. Malhame and C.-Y. Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Trans. Automat. Control*, 30(9):854 – 860, Sep 1985.
30. 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.
31. 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.
32. J. Mathieu. *Modeling, Analysis, and Control of Demand Response Resources*. PhD thesis, University of California at Berkeley, 2012.
33. J. Mathieu and D. Callaway. State estimation and control of heterogeneous thermostatically controlled loads for load following. In *45th International Conference on System Sciences*, pages 2002–2011, Hawaii, 2012. IEEE.
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. S. Meyn, P. Barooah, A. Bušić, and J. Ehren. Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in Florida. In *Proceedings of the 52nd IEEE Conf. on Decision and Control*, pages 6946–6953, Dec 2013.
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. R. Pedroncelli. Frequency Regulation Compensation in the Organized Wholesale Power Markets – FERC 755. FERC Docket Nos. RM11-7-000 and AD10-11-000; Order No. 755 – Online, <http://tinyurl.com/FERC755>, October 20 2011.
38. M. L. Puterman. *Markov Decision Processes*. Wiley, New York, 1994.
39. M. Roozbehani, M. A. Dahleh, and S. K. Mitter. Volatility of power grids under real-time pricing. *IEEE Transactions on Power Systems*, 27(4):1926–1940, 2012.
40. 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.
41. F. C. Schweppe. Power systems ‘2000’: Hierarchical control strategies. *IEEE Spectrum*, pages 42–47, July 1978.
42. J. Sharp. Electrical load disconnect device with electronic control, Dec. 2012. US Patent 8,328,110.
43. S. H. Tindemans, V. Trovato, and G. Strbac. Decentralized control of thermostatic loads for flexible demand response. *IEEE Trans. Contr. Sys. Techn.*, 23(5):1685–1700, 2015.
44. E. Todorov. Linearly-solvable Markov decision problems. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1369–1376. MIT Press, Cambridge, MA, 2007.
45. E. Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009.
46. L. C. Totu. *Large Scale Demand Response of Thermostatic Loads*. PhD thesis, Faculty of Engineering and Science, Aalborg University, 2015.
47. P. Weiss. L’hypothèse du champ moléculaire et la propriété ferromagnétique. *J. Phys. Theor. Appl.*, 6(1):661–690, 1907.
48. C. Ziras, E. Vrettos, and G. Andersson. Primary frequency control with refrigerators under startup dynamics and lockout constraints. In *IEEE Power & Energy Society General Meeting*, pages 1–5. IEEE, 2015.