

Ancillary Service to the Grid

Using Intelligent Deferrable Loads

Sean Meyn, Prabir Barooah, Ana Bušić, Yue Chen, and Jordan Ehren

Abstract

Renewable energy sources such as wind and solar power have a high degree of unpredictability and time-variation, which makes balancing demand and supply challenging. One possible way to address this challenge is to harness the inherent flexibility in demand of many types of loads. Introduced in this paper is a technique for decentralized control for automated demand response that can be used by grid operators as ancillary service for maintaining demand-supply balance.

A Markovian Decision Process (MDP) model is introduced for an individual load. A randomized control architecture is proposed, motivated by the need for decentralized decision making, and the need to avoid synchronization that can lead to large and detrimental spikes in demand. An aggregate model for a large number of loads is then developed by examining the mean field limit. A key innovation is an LTI-system approximation of the aggregate nonlinear model, with a scalar signal as the input and a measure of the aggregate demand as the output. This makes the approximation particularly convenient for control design at the grid level.

The second half of the paper contains a detailed application of these results to a network of residential pools. Simulations are provided to illustrate the accuracy of the approximations and effectiveness of the proposed control approach.

This research is supported by the NSF grant CPS-0931416, the Department of Energy Awards DE-OE0000097 & DE-SC0003879, and the French National Research Agency grant ANR-12-MONU-0019. We acknowledge the help of Mark Rosenberg who offered many suggestions to improve the manuscript, and caught several typos in earlier drafts.

S.M., Y.C. and J.E. are with the Dept. ECE and P.B. is with the Dept. MAE at the University of Florida, Gainesville. A.B. is with INRIA and the Computer Science Dept. of École Normale Supérieure, Paris, France.

I. INTRODUCTION

Renewable energy penetration is rising rapidly throughout the world, and bringing with it high volatility in energy supply. Resources are needed to compensate for these large fluctuations in power. The federal energy regulatory commission (FERC) in conjunction with generation and utility companies are struggling to find resources, and finding ways to properly compensate for ancillary services that are badly needed by each *balancing authority* (BA) in the U.S.. FERC orders 755 and 745 are examples of their attempts to provide incentives.

This paper concerns decentralized control of a large number of electric loads in a power grid. A particular load has a service it is intended to provide – clean dishes, hot water, or a clean pool. It is assumed that each load has some flexibility in energy consumption. This flexibility is harnessed to provide ancillary services to the power grid to help maintain stability, and to help offset any volatility in the grid because of line or generation outage, or because of the volatile nature of renewable energy. This is commonly called “demand response”, but the meaning is slightly different here: The tuning of energy consumption is automated, and we assume that the consumers do not suffer any degradation in the service offered by the loads.

We argue that most of the load in the U.S. is highly flexible, and this flexibility can be harnessed to provide ancillary service without central control, and without significant impact on the needs of consumers or industry. A defining characteristic of ancillary service is that on average it is a *zero-energy* service, so that the desired power consumption level to be tracked is zero on average. This makes use of deferrable loads particularly attractive as sources of ancillary service.

Many utilities already employ demand response programs that use deferrable loads to reduce peak demand and manage emergency situations. Florida Power and Light (FPL), for example, has 780,000 customers enrolled in their *OnCall Savings Program* in which residential air conditioners, water heaters, and pool pumps systems are automatically controlled when needed [1]. Today, FPL uses this service only 3–4 times per year [1]. While a valuable service to the grid, there is tremendous additional potential from these sources that today is virtually untapped.

Nearly all of America’s ISOs/RTOs also allow for demand side resources to participate in their regulation and spinning reserve markets, but as of the summer of 2013, only PJM allows aggregation (with approval) [2]. Growth of these resources in these wholesale markets has helped lower costs per megawatt-hour from 2009 to 2011 [2]. Still, markets for regulation and spinning reserves from traditional generation sources continue to grow because of increasing dependency on renewable generation.

Fig. 1 shows the regulation signal for a typical week within the Bonneville Power Authority (BPA) [3]. Its role is analogous to the control signal in the feedback loop in a flight control system. Just like in an aviation control system, the variability seen in this figure is in part a consequence of variability of wind generation in this region.

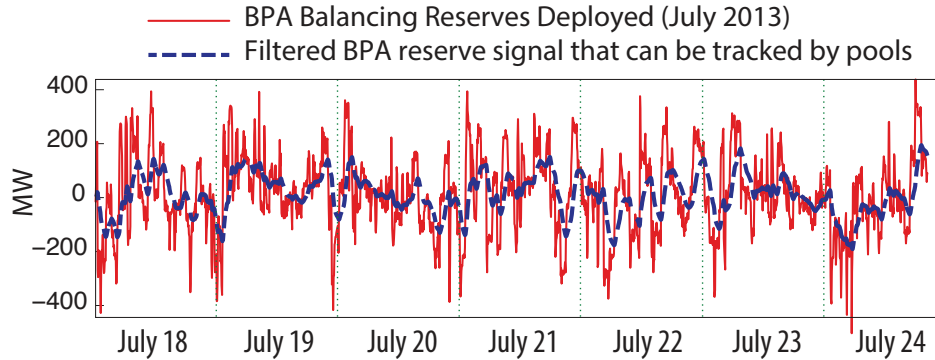


Fig. 1. *BPA Balancing Reserves Deployed* — Ancillary service needs at the BPA during one week in 2013. The maximum is approximately one-tenth of maximum load in this region.

We propose to break up a regulation signal into frequency bands for the purposes of ancillary services provisioning by various resources. In prior work it is shown how heating and ventilation systems in commercial buildings can provide service in the high frequency band, corresponding to periods ranging from 3 minutes to one hour [4]–[6]. At the lowest frequencies, an important resource will be flexible manufacturing. An example is Alcoa, that today provides 70MW of service to MISO by providing control over their aluminum smelting operation in Indiana. Alcoa’s service is provided continuously, and provides significant revenue to Alcoa and even greater benefits to the region managed by MISO.

The technical content of the paper starts with a control architecture designed to address privacy concerns and communication constraints. It is assumed that an individual load can view a regulation signal, much as we can view BPA’s regulation signal online today.

To provide ancillary service in a specified frequency band, we argue that it is essential to introduce randomization at each load. Among many benefits, randomization avoids synchronization, much like randomized congestion avoidance protocols in communication networks. First deployed nearly fifty years ago, ALOHA may be the first distributed communication protocol based on randomization. *Random Early Detection* for congestion control was introduced in the highly influential paper [7]. The historical discussion in this paper points to significant research on randomized algorithms beginning in the early 1970s, following wide deployment of ALOHA. Randomized protocols are now standard practice in

communication networks [8]. It is likely that randomized algorithms will become a standard feature of the power grid of the future.

To formulate a randomized control strategy, a Markovian Decision Process (MDP) model is proposed for an individual load. An aggregate model for a large number of loads is then obtained as a mean field limit.

A particular formulation of Todorov [9] is adopted because we can obtain an explicit solution, and because of available tools for analysis borrowed from the theory of large deviations. In particular, a key innovation in the present paper is an LTI-system approximation of the aggregate nonlinear model, which is possible through application of results from [10]. The scalar input in this linear model is a parameter that appears in the MDP cost function.

The LTI approximation is convenient for control design at the grid level: the input becomes the control signal that the BA will broadcast to all the loads, which adjusts a parameter in the randomized policy for the optimal MDP solution at each load.

In the second half of this paper we apply the general results of this paper to show how pool pumps can be harnessed to obtain ancillary service in a medium frequency band, corresponding to the dashed line in Fig. 1. This is the same BPA regulation signal, passed through a low pass filter

A pool pump is the heart of a pool's filtration system: It runs each day for a period of time range from 4 to 24 hours, and consumes over 1 KW of power when in operation [11]. The ability to control just half of Florida's pool pumps amounts to over 500 MW of power! Much of the control infrastructure is already in place [12]. Still, constraints and costs must be satisfied. These include run-times per day and per week, the cost of startup and shut down, as well as the total energy consumption. Moreover, there are privacy concerns and related communication constraints. Consequently, control algorithms must be distributed so that most of the required intelligence resides at individual pool pumps. In this paper we focus on constraints related to run-times per day, which is critical for keeping the water in the pool clean. Privacy and communication constraints will be addressed through the distributed control architecture.

A number of recent works have explored the potential for flexible loads for providing ancillary service. These include commercial building thermostatic loads to provide ancillary service in the time-scale of a few minutes (see [13] and refs. therein), electric vehicle charging [13]–[16] that can provide ancillary service in the time scale of a few hours, and our own recent work on harnessing ancillary service from commercial building HVAC [4]–[6].

Mean-field games have been employed for analysis of aggregate loads in several recent papers [14], [16]. See [17]–[19] for more on general theory of mean-field techniques.

The work of [13] is most closely related to the present paper, in that the authors also consider an aggregate model for a large collection of loads. The natural state space model is bilinear, and converted to a linear model through division of the state. The control architecture consists of a centralized control signal computation based on state feedback, and the resulting input is broadcast to the devices.

In this paper, intelligence is concentrated at the individual load: An MDP control solution is obtained at each load, but the aggregate behavior is well approximated by a *single-input single-output, linear time-invariant* (SISO-LTI) system. Hence the control problem for the balancing authority can be addressed using classical control design methods. State estimation is not required — the information required at the BA is an estimate of the proportion of loads that are operating.

In the numerical example considered in this paper, the linear system is minimum-phase and stable, which is very helpful for control design.

The remainder of the paper is organized as follows. The control solution for a single pool is described in Section II, along with approximations of the optimal control solution based on general theory presented in the Appendix. The control of the aggregate collection of pools is considered in Section III. Conclusions and directions of future research are contained in Section IV.

II. OPTIMAL CONTROL FOR A LOAD AND FOR THE GRID




Fig. 2. The control architecture: command ζ is computed at a BA, and transmitted to each pool pump. The control decision at a load is binary (turn on/off), and is based only on its own state and the signal ζ .

A. Control architecture overview

We begin with a description of the control and information architecture that is the subject of this paper. The components of the architecture are illustrated in Fig. 2:

- (i) There are N homogeneous loads that receive a common scalar command signal from the balancing authority, or BA, denoted $\zeta = \{\zeta_t\}$ in the figure.

Randomization at each load is desirable to avoid synchronization of loads, and also to facilitate analysis of the aggregate system. It is assumed that each load evolves as a controlled Markov chain: The transition probability for each load is determined by its own state, and the BA signal ζ . The common dynamics are defined by a controlled transition matrix $\{P_\zeta : \zeta \in \mathbb{R}\}$. For the i th load, there is a state process \mathbf{X}^i whose transition probability at time t is given by,

$$\mathbb{P}\{X_{t+1}^i = x^+ \mid X_t^i = x^-, \zeta_t = \zeta\} = P_\zeta(x^-, x^+) \quad (1)$$

where x^- and x^+ are possible state-values. The details of the model are described in Section II-B.

(ii) The BA has measurements of the other two scalar signals shown in the figure: The normalized aggregate power consumption \mathbf{y} and desired deviation in power consumption \mathbf{r} .

When $\zeta_t = 0$ for all t , then the aggregate power consumption takes the value \mathbf{y}^0 . The goal of the BA is a tracking problem: Achieve $y_t \approx y^0 + r_t$ for all t . This can be addressed using classical control techniques if the dynamics from ζ to $\tilde{\mathbf{y}} = \mathbf{y} - \mathbf{y}^0$ can be approximated by an LTI system.

The main contributions of this paper are based on the construction of the controlled transition matrix for an individual load, taking into account potentially conflicting goals: The BA desires overall dynamics from ζ to \mathbf{y} that facilitate tracking the reference signal \mathbf{r} . Each load requires good quality of service. In the case of a pool, the water must be kept clean, and the electricity bill must remain constant over each month.

An approach of Todorov [9] is adopted to construct the family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$. They are smooth in the parameter ζ , and a first-order Taylor series approximation gives, for any pair of states (x^-, x^+)

$$P_\zeta(x^-, x^+) = \exp(\zeta\Gamma(x^-, x^+) + O(\zeta^2))P_0(x^-, x^+) \quad (2)$$

where P_0 denotes the dynamics of a load when $\zeta \equiv 0$, and Γ is a matrix. Based on (27), we have

$$\Gamma(x^-, x^+) = \tilde{U}(x^-) + H(x^+) - H(x^-)$$

where the function H is a solution to *Poisson's equation*; a linear equation for the nominal model.

This structure leads to the LTI approximation of the input-output dynamics from ζ to $\tilde{\mathbf{y}}$ that is presented in Proposition 2.4. Section II-C also contains second-order approximations of P_ζ .

In Section III these general techniques are applied to a collection of residential pools. In this example it is found that the LTI model is minimum phase, and that a simple PI controller can be effectively used for the control transfer function G_c shown in Fig. 2.

B. Load model and design

In this section we present a procedure to construct the controlled transition matrix appearing in (1). The controlled Markov chain evolves on a finite state space, denoted $\mathsf{X} = \{x^1, \dots, x^d\}$. The construction is based on an optimal control problem for an individual load, taking into account the needs of the load and the grid.

It is assumed that a transition matrix P_0 is given that models “control free” behavior of the Markov chain, and a utility function $\mathcal{U}: \mathsf{X} \rightarrow \mathbb{R}$ is used to model the needs of the grid. The optimal control problem will balance average utility and the cost of deviation.

Since we focus on a single load, in this subsection the index i in (1) is dropped, and we denote by $\mathbf{X} = (X_0, X_1, \dots)$ the stochastic process evolving on X that models this load.

In the second half of the paper we will focus on a particular example in which each load is a residential pool pump. The true nominal behavior would be deterministic – most consumers set the pump to run a fixed number of hours each day. However, the randomized policy is based on a stochastic model for nominal behavior, so we introduce some randomness to define the nominal transition matrix P_0 . The state space is taken to be the finite set,

$$\mathsf{X} = \{(m, i) : m \in \{\oplus, \ominus\}, i \in \{1, \dots, T\}\} \quad (3)$$

If $X_t = (\ominus, i)$, this indicates that the pool-pump was turned off and has remained off for i time units, and $X_t = (\oplus, i)$ represents the alternative that the pool-pump has been operating continuously for exactly i time units. A state-transition diagram is shown in Fig. 3. The values of $P_0(x, y)$ will be chosen to be nearly 0 or 1 for most $x, y \in \mathsf{X}$.




Fig. 3. State transition diagram for the pool-pump model.

The utility function \mathcal{U} on X is chosen as the indicator function that the pool pump is operating:

$$\mathcal{U}(x) = \sum_i \mathbb{I}\{x = (\oplus, i)\}$$

Whether this actually represents any utility to the grid operator depends on the state of the grid. This will be clarified after we define the optimization problem and its solution.

We now return to the general model. Consider first a finite-time-horizon optimization problem: For a given terminal time T , let p_0 denote the probability measure on strings of length T :

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

where $x_0 \in \mathbf{X}$ is assumed to be given.

A fixed scalar $\zeta \in \mathbb{R}$ is interpreted as a *weighting parameter* in the following definition of *total welfare*. For any probability measure p , this is defined as the weighted difference,

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

where the expectation is with respect to p , and D denotes relative entropy. Let p^{T*} denote the probability measure that maximizes this expression.

Proposition 2.1: The probability measure p^{T*} is the *twisted distribution*,

$$p^{T*}(x_1, \dots, x_T) = \exp \left(\zeta \sum_{t=1}^T \mathcal{U}(x_t) - \Lambda_T(\zeta) \right) p_0(x_1, \dots, x_T) \quad (4)$$

where

$$\Lambda_T(\zeta) = \log \left\{ \mathbb{E} \left[\exp \left(\zeta \sum_{t=1}^T \mathcal{U}(X_t) \right) \right] \right\}, \quad (5)$$

and the expectation is with respect to p_0 . Moreover, $\mathcal{W}_T(p^{T*}) = \Lambda_T(\zeta)$ is the optimal value. \square

Proof: Optimality of p^{T*} follows from convex duality between the log moment generating function and relative entropy – [20, Proposition II.1] and [21, Lemma 2.39]. The formula (5) follows from the fact that p^{T*} sums to unity, so that $\Lambda_T(\zeta)$ can be interpreted as a normalizing constant.

The identity $\mathcal{W}_T(p^{T*}) = \Lambda_T(\zeta)$ follows from the definitions of \mathcal{W}_T and p^{T*} . \square

The probability measure p^{T*} defines a Markov chain on the time interval $\{0, 1, \dots, T\}$, but it is not necessarily time-homogeneous. In the infinite horizon case, we would like to find a distribution p^* on infinite sequences that attains the optimal average welfare,

$$\eta_\zeta^* = \lim_{T \rightarrow \infty} \frac{1}{T} \mathcal{W}_T(p^{T*}) = \lim_{T \rightarrow \infty} \frac{1}{T} \log \left\{ \mathbb{E} \left[\exp \left(\zeta \sum_{t=1}^T \mathcal{U}(X_t) \right) \right] \right\} \quad (6)$$

A solution to the infinite horizon problem is given by a time-homogenous Markov chain whose transition matrix \tilde{P}_ζ is easy to compute, based on the solution of an eigenvector problem; these results are summarized in the proposition that follows. The proof of Proposition 2.2 is given in the Appendix.

Proposition 2.2: If P_0 is irreducible, an optimizing p^* that achieves (6) is defined by a time-homogeneous Markov chain whose transition probability is given by

$$\tilde{P}_\zeta(x, y) = \frac{1}{\lambda} \frac{1}{v(x)} \hat{P}_\zeta(x, y) v(y), \quad x, y \in \mathsf{X}, \quad (7)$$

where \hat{P}_ζ is the scaled transition matrix,

$$\hat{P}_\zeta(x, y) = \exp(\zeta \mathcal{U}(x)) P_0(x, y), \quad x, y \in \mathsf{X}, \quad (8)$$

and λ, v is that eigen-pair corresponding to the eigenvector problem

$$\hat{P}_\zeta v = \lambda v \quad (9)$$

such that $\lambda = \lambda_\zeta > 0$ is the unique maximal eigenvalue for \hat{P}_ζ , $v = v_\zeta$ is unique up to constant multiples, and $v(x) > 0$ for each $x \in \mathsf{X}$. In addition, the following bounds hold for each T ,

$$\begin{aligned} 0 \leq \mathcal{W}_T(p^{T*}) - \mathcal{W}_T(\tilde{p}^T) &\leq 2\|h\|_{\text{sp}} \\ |T\Lambda - \mathcal{W}_T(p^{T*})| &\leq \|h\|_{\text{sp}} \end{aligned} \quad (10)$$

where $h = \log v$, and the span norm is defined by $\|h\|_{\text{sp}} = \max h - \min h$.

Consequently, the Markov model achieves the optimal average welfare (6) with $\eta_\zeta^* = \Lambda$. \square

The eigenvector problem 9 appears in multiplicative ergodic theory [10], and also in Todorov's analysis [9]. It is shown in [9] that the *relative value function* appearing in the average cost optimality equations is the logarithm of the eigenvector:

$$h^*(x) = \log(v(x)), \quad x \in \mathsf{X}. \quad (11)$$

See also the derivation in [22] for a variant of this model.

Second order Taylor series approximations for v and η^* near $\zeta \approx 0$ can be found by borrowing tools from large-deviations theory. Some of these approximation results are new, and are collected together in the next section and in the Appendix.

C. Approximations

Approximations will be needed for analysis when we extend the model to allow ζ to change with time.

A solution to the eigenvector problem (9) can be represented through a regenerative formula. Let $\alpha \in \mathsf{X}$ be some fixed state that is reachable from each initial condition of the chain, under the transition law P_0 . That is, the chain is assumed to be α -irreducible [23]. Since the state space is assumed to be finite, it follows that there is a unique invariant probability measure π_0 for P_0 . The first return time is denoted,

$$\tau = \min\{t \geq 1 : X_t = \alpha\}.$$

Recall that the infinite horizon optimal welfare is given by $\eta_\zeta^* = \log(\lambda)$. From the theory of positive matrices [10], [24], [25], it follows that it is the unique solution to,

$$1 = \mathbf{E}_\alpha \left[\exp \left(\sum_0^{\tau-1} [\zeta \mathcal{U}(X_t) - \eta_\zeta^*] \right) \right] \quad (12)$$

where the subscript indicates that the initial condition is $X(0) = \alpha$. Moreover, for each $x \in \mathsf{X}$, the value of $v(x)$ is obtained as the expected sum, with initial condition $X(0) = x$:

$$v(x) = \mathbf{E}_x \left[\exp \left(\sum_0^{\tau-1} [\zeta \mathcal{U}(X_t) - \eta_\zeta^*] \right) \right] \quad (13)$$

These expectations are each with respect to the nominal transition law P_0 .

A Taylor-series approximation of η_ζ^* is based on two parameters, defined with respect to the nominal model P_0 with invariant probability measure π_0 . The first-order coefficient is the the steady-state mean of \mathcal{U} ,

$$\eta_0 = \sum_x \pi_0(x) \mathcal{U}(x) \quad (14)$$

The second-order coefficient is based on the *asymptotic variance* of \mathcal{U} for the nominal model (the variance appearing in the Central Limit Theorem (CLT) for the nominal model). For this finite state space model this has two similar representations,

$$\begin{aligned} \kappa^2 &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\left(\sum_0^{T-1} \tilde{\mathcal{U}}(X_t) \right)^2 \right] \\ &= \pi_0(\alpha) \mathbf{E}_\alpha \left[\left(\sum_0^{\tau-1} \tilde{\mathcal{U}}(X_t) \right)^2 \right] \end{aligned} \quad (15)$$

where $\tilde{\mathcal{U}} = \mathcal{U} - \eta_0$. See [23, Theorem 17.0.1] for the CLT, and eqn. (17.13) of [23] for the second representation above.

Similarly, the following functions of x are used to define a second order Taylor series approximation for h_ζ^* . The first-order term is the solution to *Poisson's equation* for P_0 ,

$$H(x) = \mathbf{E}_x \left[\sum_0^{\tau-1} \tilde{\mathcal{U}}(X_t) \right] \quad (16)$$

The asymptotic variance can be expressed in terms of Poisson's equation [23], [26]:

$$\kappa^2 = \sum_x \pi_0(x) (2\tilde{\mathcal{U}}(x)H(x) - \tilde{\mathcal{U}}(x)^2)$$

The second-order term in an approximation of v is another variance,

$$\mathcal{S}(x) = \mathbf{E}_x \left[\left(\sum_0^{\tau-1} \tilde{\mathcal{U}}(X_t) \right)^2 \right] - (H(x))^2, \quad x \in \mathsf{X}. \quad (17)$$

Proposition 2.3: The following hold for the finite state space model in which P_0 is irreducible:

(i) The optimal average welfare η_ζ^* is convex as a function of ζ , and admits the Taylor series expansion,

$$\eta_\zeta^* = \eta_0\zeta + \frac{1}{2}\kappa^2\zeta^2 + O(\zeta^3) \quad (18)$$

(ii) The mean of \mathcal{U} under the the invariant probability measure $\tilde{\pi}_\zeta$ for \check{P}_ζ is given by,

$$\sum_x \tilde{\pi}_\zeta(x)\mathcal{U}(x) = \frac{d}{d\zeta}\eta_\zeta^* \quad (19)$$

This admits the first-order Taylor series approximation

$$\frac{d}{d\zeta}\eta_\zeta^* = \eta_0 + \kappa^2\zeta + O(\zeta^2) \quad (20)$$

(iii) The relative value function (11) admits the second-order Taylor series approximation,

$$h_\zeta^*(x) = \zeta H(x) + \frac{1}{2}\zeta^2 \mathcal{S}(x) + O(\zeta^3) \quad (21)$$

□

Proof: Equations (18)—(20) follow from the fact that $\eta_\zeta^* = \log(\lambda)$ can be expressed as a cumulative log-moment generating function [10, Prop. 4.9].

Convexity follows from the fact that η_ζ^* is the maximum of linear functions of ζ (following the linear-program formulation of the ACOE [27]).

The approximation (21) follows from the the representation (13) for v , and the definition $h_\zeta^* = \log(v)$ (see (11)). □

The representations in this subsection are useful for analysis, but not for computation. Methods to compute H and \mathcal{S} are contained in the Appendix.

D. Aggregate load model

Consider N loads operating independently under the randomized policy described in the previous section. The state of the i th load is denoted X_t^i . For large N we have from the Law of Large Numbers,

$$\frac{1}{N} \sum_{i=1}^N \mathcal{U}(X_t^i) \approx \mathbb{E}[\mathcal{U}(X_t)] \quad (22)$$

The expectation and probability on the right are with respect to the optimal transition law \check{P}_ζ , where ζ is the parameter used in (6).

We pose the following centralized control problem: How to choose the variable ζ to regulate average utility *in real time*, based on measurements of the average utility, and also a regulation signal denoted r . Let y_t be the fraction of loads that are on at time t :

$$y_t = \frac{1}{N} \sum_{i=1}^N \mathcal{U}(X_t^i), \quad (23)$$

which is assumed to be observed by the BA.

To address the control problem faced by the BA it is necessary to relax the assumption that this parameter is fixed. We let $\zeta = \{\zeta_0, \zeta_1, \dots\}$ denote a sequence of scalars, which is regarded as an input signal for the control problem faced by the BA. An aggregate model is obtained in two steps.

In step 1 the existence of a mean-field limit is assumed: Let $N \rightarrow \infty$ to obtain the generalization of (22),

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{X_t^i = x\} = \mu_t(x), \quad x \in \mathcal{X}. \quad (24)$$

For a given initial distribution μ_0 on \mathcal{X} , the distribution μ_t is defined by $\mu_t(x_t) =$

$$\sum_{x_i \in \mathcal{X}} \mu_0(x_0) \check{P}_{\zeta_0}(x_0, x_1) \check{P}_{\zeta_1}(x_1, x_2) \cdots \check{P}_{\zeta_{t-1}}(x_{t-1}, x_t) \quad (25)$$

where x_t is an arbitrary state in \mathcal{X} , and the sum is over all intermediate states. We view $\{\mu_t\}$ as a state process that is under our control through ζ .

Justification for the mean-field limit is contained in Theorem 2.5.

Step 2 is based on the Taylor series approximations surveyed in the previous section to approximate this nonlinear system by a linear state space model with d -dimensional state Φ and output γ . It is defined so that for any time t , and any i ,

$$\mu_t(x^i) = \pi_0(x^i) + \Phi_t(i) + o(\zeta)$$

$$\gamma_t = \tilde{y}_t + o(\zeta)$$

where $\tilde{y}_t = y_t - y^0$, with $y^0 = \sum_x \pi_0(x) \mathcal{U}(x)$, and where $o(\zeta)$ is in fact $O(\zeta_0^2 + \dots + \zeta_t^2)$.

Proposition 2.4: Consider the nonlinear state space model whose state evolution is $\mu_{t+1} = \mu_t \check{P}_{\zeta_t}$, and output is $y_t = \sum_x \mu_t(x) \mathcal{U}(x)$. Its unique equilibrium with $\zeta \equiv 0$ is $\mu_t \equiv \pi_0$ and $y_t \equiv y^0 := \sum_x \pi_0(x) \mathcal{U}(x)$. Its linearization around this equilibrium is given by,

$$\begin{aligned} \Phi_{t+1} &= A\Phi_t + B\zeta_t \\ \gamma_t &= C\Phi_t \end{aligned} \quad (26)$$

where $A = P_0^T$, C is a row vector of dimension $d = |\mathsf{X}|$ with $C_i = \mathcal{U}(x^i)$ for each i , and B is a d -dimensional column vector with entries $B_j = \sum_x \pi_0(x) \mathcal{E}(x, x^j)$, where

$$\mathcal{E}(x^i, x^j) = \left[\tilde{\mathcal{U}}(x^i) + H(x^j) - H(x^i) \right] P_0(x^i, x^j) \quad (27)$$

for each $x^i, x^j \in \mathsf{X}$. The initial condition is $\Phi_0(i) = \mu_0(x^i) - \pi_0(x^i)$, $1 \leq i \leq d$.

The matrix \mathcal{E} is equal to the derivative,

$$\mathcal{E} = \left. \frac{d}{d\zeta} P_\zeta \right|_{\zeta=0}$$

Consequently, the formula (27) implies the approximation (2). \square

Proof: The formulae for A and C follow from the fact that the system is linear in the state. We have, from (7),

$$\check{P}_\zeta(x^i, x^j) = e^{\zeta \mathcal{U}(x^i) - \eta_\zeta^* - h_\zeta^*(x^i)} P_0(x^i, x^j) e^{h_\zeta^*(x^j)}$$

Based on the first order approximation of h_ζ^* in Proposition 2.3 we obtain,

$$\check{P}_\zeta(x^i, x^j) \approx e^{\zeta[-H(x^i) + \tilde{\mathcal{U}}(x^i)]} P_0(x^i, x^j) e^{\zeta H(x^j)}$$

where H is a solution to Poisson's equation (with forcing function \mathcal{U}) for the nominal model (see (16)).

Using a first order Taylor series for the exponential then gives,

$$\begin{aligned} \check{P}_\zeta(x^i, x^j) &\approx [1 - \zeta(H(x^i) - \tilde{\mathcal{U}}(x^i))] P_0(x^i, x^j) [1 + \zeta H(x^j)] \\ &\approx P_0(x^i, x^j) + \zeta \mathcal{E}(x^i, x^j) \end{aligned}$$

If $\mu \approx \pi_0$ and ζ is small, then we can approximate,

$$\mu \check{P}_\zeta \approx \mu P_0 + \zeta B^T,$$

where B is the column vector with entries $B_j = \sum_x \pi_0(x) \mathcal{E}(x, x^j)$. \square

Next we justify the mean-field model (24).

For the purpose of analysis we lift the state space from the d -element set $\mathsf{X} = \{x^1, \dots, x^d\}$, to the d -dimensional simplex \mathbf{S} . For the i^{th} load at time t , the element $\pi_t^i \in \mathbf{S}$ is the degenerate distribution whose mass is concentrated at x if $X_t^i = x$. The average over N , denoted $\mu_t^N \in \mathbf{S}$, is the empirical distribution,

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

In the proof of convergence it is assumed that ζ^N is obtained using state feedback of the form,

$$\zeta_t^N = \phi_t(\mu_0^N, \dots, \mu_t^N)$$

where $\phi_t: \mathbf{S}^{t+1} \rightarrow \mathbb{R}$ is continuous for each t , and does not depend upon N . The following result establishes convergence.

Theorem 2.5: Suppose $\mu_0^N \rightarrow \mu_0$ as $N \rightarrow \infty$, and that the state transition matrix P_ζ is continuous as a function of ζ . Then for each t ,

$$\lim_{N \rightarrow \infty} \mu_t^N = \mu_t, \quad \text{with probability one,} \quad (28)$$

where the right hand side denotes the probability measure (25), in which

$$\zeta_t = \phi_t(\mu_0, \dots, \mu_t), \quad t \geq 0.$$

□

The proof of this result is given at the end of this subsection, and is largely based on a version of the Law of Large Numbers.

Let $\{M_{N,k}, 1 \leq k \leq N\}$ denote a martingale array: This means that $E[M_{N,k} | M_{N,1}, \dots, M_{N,j}] = M_{N,j}$ for each N and $1 \leq j < k \leq N$. When $k = N$, we denote $M_N = M_{N,N}$.

Proposition 2.6: Suppose that $M_{N,k}$ is a martingale array with bounded increments: For some $c_m < \infty$,

$$|M_{N,k+1} - M_{N,k}| \leq c_m \quad \text{for all } k \text{ and } N$$

Then the Law of Large Numbers holds:

$$\lim_{N \rightarrow \infty} \frac{M_N}{N} = 0, \quad \text{with probability one.}$$

□

Proof: The Hoeffding-Azuma inequality [28] gives the following bound:

$$\mathbf{P}\{N^{-1}|M_N| \geq t\} \leq 2 \exp(-[Nt]^2/[2Nc_m^2])$$

The right hand is summable, so the result follows from the Borel-Cantelli Lemma. □

Proposition 2.6 is applied to show that the sequence of empirical distribution μ_t^N can be approximated by the mean-field model perturbed by a disturbance that vanishes as $N \rightarrow \infty$:

Lemma 2.7: The empirical distributions $\{\mu_t^N : t \geq 0\}$ obey the recursion

$$\mu_{t+1}^N = \mu_t^N P_{\zeta_t^N} + W_{t+1}^N, \quad (29)$$

in which, $W_{t+1}^N = \frac{1}{N} \sum_{i=1}^N \Delta_{t+1}^i$ for a family of vector random variables $\{\Delta_{t+1}^i\}$.

On denoting $M_{N,k} = \sum_{i=1}^k \Delta_t^i$ we have,

(i) $\{M_{N,k} : 1 \leq k \leq N\}$ is a martingale array.

(ii) There exists c_m such that $\|M_{N,k} - M_{N,k-1}\| \leq c_m$ for all N and all k such that $1 < k \leq N$.

Proof of Lemma 2.7: To establish (29) we first establish a similar expression for $\{\pi_t^i\}$.

For each i , the sequence of degenerate distributions $\{\pi_t^i\}$ evolve according to a random linear system,

$$\pi_{t+1}^i = \pi_t^i G_{t+1}^i \quad (30)$$

in which π_t^i is interpreted as a d -dimensional row vector, and G_t^i is a $d \times d$ matrix with entries 0 or 1 only, and $\sum_l G_t^i(x^j, x^l) = 1$ for all j . It is conditionally independent of $\{\pi_0^i, \dots, \pi_t^i\}$, given ζ_t^N , with

$$\mathbb{E}[G_{t+1}^i | \pi_0^i, \dots, \pi_t^i, \zeta_t^N] = P_{\zeta_t^N}. \quad (31)$$

Dependency of π_t^i, G_t^i on N is suppressed, but we must distinguish ζ_t^N from its limit ζ_t .

The random linear system (30) can thus be described as a linear system driven by “white noise”:

$$\pi_{t+1}^i = \pi_t^i P_{\zeta_t^N} + \Delta_{t+1}^i \quad (32)$$

where, $\{\Delta_{t+1}^i = \pi_t^i(G_{t+1}^i - P_{\zeta_t^N}) : t \geq 1\}$, which establishes (29).

The following representation will clarify the remaining analysis:

$$G_t^i = \mathcal{G}(\zeta_{t-1}^N, \xi_t^i), \text{ where } \{\xi_t^i : t \geq 1, i \geq 1\} \text{ are i.i.d.} \quad (33)$$

For $1 \leq i < N$ and fixed t , we define two σ -algebras:

$$\mathcal{F}_i = \sigma\{\Delta_t^k, k \leq i\}$$

$$\mathcal{H}_i = \sigma\{\pi_{t-1}^{k+1}, \zeta_{t-1}^N, \Delta_t^k, k \leq i\}$$

Under (33) we have the extension of (31), that $\mathbb{E}[G_t^{i+1} | \mathcal{H}_i] = P_{\zeta_{t-1}^N}$. Moreover, by construction the random variable π_{t-1}^{i+1} is \mathcal{H}_i -measurable. Therefore,

$$\mathbb{E}[\Delta_t^{i+1} | \mathcal{H}_i] = \mathbb{E}[\pi_{t-1}^{i+1}(G_t^{i+1} - P_{\zeta_{t-1}^N}) | \mathcal{H}_i] = 0$$

The smoothing property of the conditional expectation, and the construction $\mathcal{F}_i \subset \mathcal{H}_i$, then gives (i),

$$\mathbb{E}[\Delta_t^{i+1} | \mathcal{F}_i] = \mathbb{E}[\mathbb{E}[\Delta_t^{i+1} | \mathcal{H}_i] | \mathcal{F}_i] = 0$$

From the definition of Δ_t^i below equation (32), it follows that $\{\|\Delta_t^i\|\}$ admits a uniform bound. Consequently, $\|M_{N,k} - M_{N,k-1}\| = \|\Delta_t^k\|$ is bounded, which is (ii). \square

Proof of Theorem 2.5: Denote, for $T \geq 0$, the deviation $\tilde{\mu}_T^N = \mu_T^N - \mu_T$.

We prove by induction on T that $\tilde{\mu}_T^N \rightarrow 0$ as $N \rightarrow \infty$. This holds by assumption when $T = 0$.

Suppose now that (28) holds for $t \leq T$. By continuity of ϕ_t , it follows that $\zeta_t^N \rightarrow \zeta_t$ as $N \rightarrow \infty$. We also have by the definitions,

$$\tilde{\mu}_{T+1}^N = \tilde{\mu}_T^N P_{\zeta_T} + \mu_T^N (P_{\zeta_T^N} - P_{\zeta_T}) + W_{T+1}^N$$

Lemma 2.7 and Proposition 2.6 imply that $W_{T+1}^N \rightarrow 0$ as $N \rightarrow \infty$. Continuity of P_ζ then implies that

$$\lim_{N \rightarrow \infty} \tilde{\mu}_{T+1}^N = 0$$

□

III. CONTROLLING A LARGE NUMBER OF POOLS

For the remainder of the paper we apply the results of the previous section to the control of a large population of residential pools. The nominal transition matrix P_0 is defined by the probabilities of turning the pump on or off, as illustrated in the state transition diagram Fig. 3. In many of the numerical results described below a symmetric model was chosen for P_0 in which $p_i^\oplus = p_i^\ominus$, where $p_i^\oplus := \mathbb{P}\{\text{pump switches on} \mid \text{it has been off } i \text{ hours}\}$. Similarly, $p_i^\ominus := \mathbb{P}\{\text{pump switches off} \mid \text{it has been on } i \text{ hours}\}$.

The utility function \mathcal{U} on \mathbb{X} is chosen as the indicator function that the pool pump is operating:

$$\mathcal{U}(x) = \sum_i \mathbb{I}\{x = (\oplus, i)\} \quad (34)$$

The parameter ζ in (6) can be positive or negative; If $\zeta > 0$ this control formulation is designed to provide incentive to turn pumps on.

It remains to give numerical values for p_i^\oplus and p_i^\ominus , $1 \leq i \leq T$. In the symmetric model, the specification of these probabilities is performed as follows. Fix $\gamma > 1$ and define,

$$\varrho_s(x) = \begin{cases} 2^{\gamma-1} x^\gamma & 0 \leq x \leq 1/2 \\ 1 - 2^{\gamma-1} (1-x)^\gamma & 1/2 \leq x \leq 1 \end{cases}$$

If over a 24 hour day we choose a sampling time $T = 30$ minutes, then in the symmetric model we take,

$$p_i^\oplus = p_i^\ominus = \varrho_s(i/48), \quad 1 \leq i \leq 48. \quad (35)$$

Fig. 4 shows a plot of the resulting probability p_i^\oplus vs. i with $\gamma = 6$.

To go beyond the asymmetric model, introduce a parameter α intended to represent the fraction of the day that the pool is operating. We modify ϱ_s as follows,

$$\varrho_s^+(x) = \varrho_s(x^{\delta^+}), \quad \varrho_s^-(x) = \varrho_s(x^{\delta^-}),$$

where δ_+ is chosen so that $0.5^{\delta_+} = 1 - \alpha$, or $\delta_+ = -\log_2(1 - \alpha)$, and similarly $\delta_- = -\log_2(\alpha)$. For the same sampling parameters as in the previous example, we then take,

$$p_i^\oplus = \varrho_s^+(i/48), \quad p_i^\ominus = \varrho_s^-(i/48), \quad 1 \leq i \leq 48. \quad (36)$$

As $\gamma \rightarrow \infty$, the functions in (36) will converge to step functions corresponding to a deterministic cleaning period of $\alpha \times 24$ hours. We find numerically that the average cleaning period is somewhat smaller when $\alpha < \frac{1}{2}$ and $\gamma < \infty$.

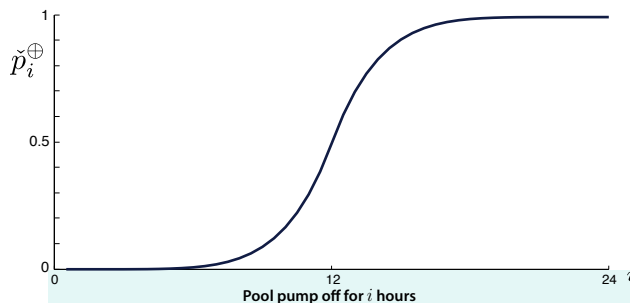


Fig. 4. Control free behavior of a pool used for numerical studies.

A. Approximations

The steady-state probability that a pool-pump is in operation is given by

$$\check{P}\{\text{pool-pump is on}\} = \sum_x \check{\pi}_\zeta(x) \mathcal{U}(x)$$

A linear approximation is obtained in Proposition 2.3 (ii):

$$\check{P}\{\text{pool-pump is on}\} = \eta_0 + \kappa^2 \zeta + O(\zeta^2) \quad (37)$$

A comparison of the true probability and its affine approximation is shown in Fig. 5 for the symmetric model, in which $\eta_0 = 1/2$. The approximation is very tight for $|z| \leq 3$. For larger values of ζ the true steady-state probability saturates (approximately 0.9 as $\zeta \rightarrow +\infty$).

For fixed ζ , the controlled model \check{P} has the same form as P_0 , with transformed probability vectors \check{p}^\oplus and \check{p}^\ominus . Fig. 6 contains plots of the transformed vector \check{p}^\oplus for values $\zeta = 0, \pm 2, \pm 4$. The plots of \check{p}^\ominus are obtained through symmetry.

The approximation of the average welfare established in Proposition 2.3 is,

$$\eta_\zeta^* = \eta_0 \zeta + \frac{1}{2} \kappa^2 \zeta^2 + O(\zeta^3) \quad (38)$$

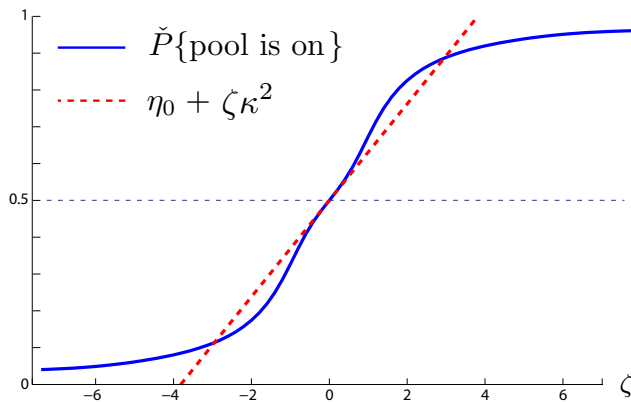


Fig. 5. Approximation of the steady-state probability that a pool-pump is operating under \tilde{P} .




Fig. 6. Transformed probability vector \tilde{p}^\oplus under \tilde{P} .

Shown in Fig. 7 is a comparison of η_ζ^* with linear and quadratic approximations based on (38).

The plots in Fig. 8 compare the eigenvector $v = e^{h\zeta}$ with the exponential of the quadratic approximation (21) given in Proposition 2.3 (iii). The computations of H and \mathcal{S} were based on the alternate expression for these functions that are described in Proposition A.1. They are normalized so that the common maxima are equal to unity. The approximation is nearly perfect for the range of $\zeta \in [-4, 4]$.

B. Aggregate load model for pool population

Here we examine the linear model (26) that will be used by the BA for control synthesis.

We begin with an equilibrium analysis in which ζ is held constant: Suppose that ζ does not vary with time, $\zeta_t = \zeta^*$ for all t , and consider the steady-state behavior of the mean-field model. We denote $y_\infty = \lim_{t \rightarrow \infty} y_t$, which is the steady-state probability that a pool is on, for the model with transition




Fig. 7. The optimal average welfare η_ζ^* and its quadratic approximation.

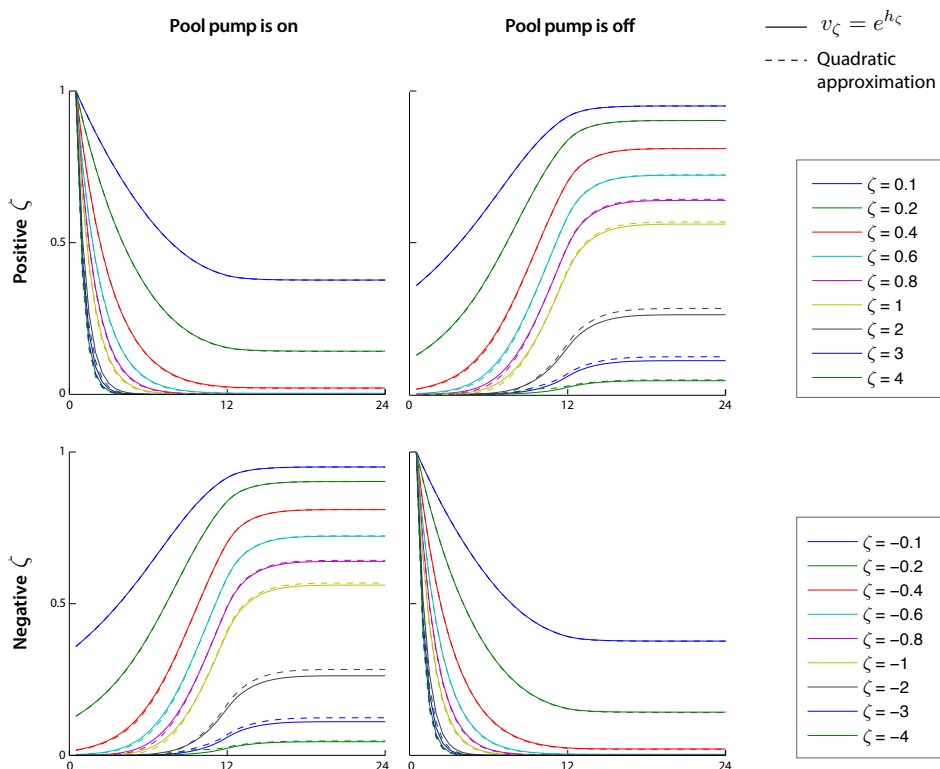


Fig. 8. Eigenvectors $v_\zeta = e^{h_\zeta}$, and their quadratic approximations $\exp(\zeta H(x) + \frac{1}{2}\zeta^2 \mathcal{S}(x))$.

law \check{P}_{ζ^*} . This can be approximated using Proposition 2.3:

$$y_\infty = \mathbb{P}\{\text{Pump is operating}\} \approx \eta_0 + \kappa^2 \zeta^*$$

From the viewpoint of the BA, there is a value G^* of desired consumption by all the pools. If $g_p > 0$ denotes the consumption of one pool pump in operation, and if there are N pools in total, then the




Fig. 9. Frequency response and pole-zero plot for the linearized model $C[Iz - A]^{-1}B$

desired steady-state probability is $y_\infty = G^*/(Ng_p)$. This translates to a corresponding value of ζ^* ,

$$\zeta^* \approx \frac{1}{\kappa^2} \left[\frac{1}{g_p} \frac{G^*}{N} - \eta_0 \right] = \frac{1}{\kappa^2} \frac{1}{g_p} \tilde{G} \quad (39)$$

where $\tilde{G} = G^* - G_0$, with $G_0 = g_p N \eta_0$, the control-free value obtained with $\zeta^* = 0$.

Consider now the case in which ζ is a function of time. Fig. 9 shows the Bode plot and pole-zero plot for the linear model (26). The transfer function from ζ to γ is BIBO stable and minimum phase.

C. Super-sampling

Recall the control architecture described at the start of Section III. At any given time, the desired power consumption/curtailment is determined by the BA based on its knowledge of dispatchable and uncontrollable generation as well as prediction of load. This is passed through a band-pass filter and scaled appropriately based on the proportion of ancillary service provided by the pools, and the average power consumption of pool pumps. The resulting reference signal is denoted r .

We introduce here a refinement of the randomized control scheme to account for delay in the system: Even if sampling takes place each hour, if a percentage of pools turn off in response to a regulation signal, then the power consumption in the grid will drop nearly instantaneously. Nevertheless, the control system model will have a one hour delay, which is unacceptable.

To obtain a more responsive system we employ “super-sampling” at the grid level, which is obtained as follows: We maintain the assumption that each pool checks the regulation signal at intervals of length T . However, the pools have no common clock.

It is convenient to model super-sampling via binning of time, so that we retain a discrete time model. Let $m > 1$ denote a “super-sampling” parameter. At the grid-level the system is in discrete time, with sampling interval T/m . For example, if $T = 30$ minutes, then $m = 6$ corresponds to a five minute sampling interval. A pool is class i if the reference signal is checked at times $nT + (i - 1)T/m$, with $n \geq 0, 1 \leq i \leq m$.

Letting y_t^i denote the fraction of pools in the i th class that are operating, the total that are operating at time t is the sum,

$$y_t = \sum_{i=0}^{m-1} y_t^i$$

Let H_0 denote the discrete time transfer function using $m = 1$, which is simply the transfer function for the linear state space model (26). For general m , the transfer function from ζ to y is

$$H(z^{-1}) = z^m H^0(z^{-m})L(z) \quad (40)$$

where L is the low pass filter,

$$L(z) = \frac{1}{m} \sum_{i=1}^m z^{-i} = \frac{1}{m} z^{-1} \frac{1 - z^{-m}}{1 - z^{-1}}$$

The term “ $1/m$ ” appears because the pools in each bin contribute this fraction of total ancillary service. In the second representation there is a pole-zero cancellation at $z = 1$. The filter $L(z)$ has $m - 1$ zeros on the unit circle: All of the solutions to $z^m = 1$, except for the solution $z = 1$.

Using super-sampling we have achieved our goal of reducing delay: In real time, the delay in this model is T/m rather than T .

D. Simulation results

The numerical results described here are based on a stochastic simulation of one million pools ($N = 10^6$), using Matlab. This large number of pools is consistent with Florida or California.

For the purposes of translation to megawatts, it is assumed that each pool in operation consumes $g_p = 1$ KW. Power consumption at time t is assumed to be equal to $N g_p y_t$ (in KW).

The supersampling approach was used in all of these experiments, with the following values of T and m fixed throughout: Each pool checks the regulation signal every $T = 30$ minutes. The supersampling parameter is $m = 12$, corresponding to 150 second sampling intervals at the grid level.

The reference signal was chosen to be the BPA regulation signal passed through a low pass filter, shown in Fig. 1. It was found that one million pools could provide far more regulation than the ± 200 MW

required at BPA during this week. More experiments were conducted in which the signal was scaled to investigate the limits of regulation from a population of one million pools.

We summarize results obtained from two sets of experiments conducted in two scenarios. In the first, the symmetric model based on the nominal model was used, with switching probability (35).

The second scenario was based on a shorter cleaning schedule of 8 hours per day. The switching probabilities defined in (36) were used in which $\alpha = 1/3$, and $\gamma = 6$ in both scenarios. The function p^\oplus using $\alpha = 1/3$ is shown in Fig. 10.




Fig. 10. Nominal model with 8 hour cleaning schedule.

In both scenarios, the linearization (26) is minimum phase: All zeros of $H_0(z) = C(Iz - A)^{-1}B$ lie strictly within the unit disk in the complex plane. With the introduction of super-sampling, the resulting transfer function (40) also has zeros on the unit circle.

In these experiments it was assumed that the BA had perfect measurements of the total power consumption of the population of pools. PI control was used to obtain the signal ζ : A proportional gain of 20, and integral gain of 4 worked well in all cases. That is, the command ζ was taken to be

$$\zeta_t = 20e_t + 4e_t^I, \quad \text{with } e_t = r_t - y_t \text{ and } e_t^I = \sum_{k=0}^t e_k$$

This is of the form $\zeta_t = \phi_t(\mu_0, \dots, \mu_t)$, $t \geq 0$, that is required in Theorem 2.5.

The average proportion of time that a pool is on will be approximately $1/2$ in Scenario 1, and $1/3$ in Scenario 2. Consequently, the class of regulation signals that can be tracked is not symmetric in Scenario 2: The population of pools has more potential for increasing rather than decreasing power consumption.

To attempt to quantify this effect, define *potential capacity* as the upper and lower limits of power deviation, subject to the constraint that tracking performance does not degrade, denoted $\{+Demand, -Supply\}$. Through simulations it was found that the potential capacity in Scenario 1 is $\{+500MW, -500MW\}$, and $\{+695MW, -305MW\}$ in Scenario 2.

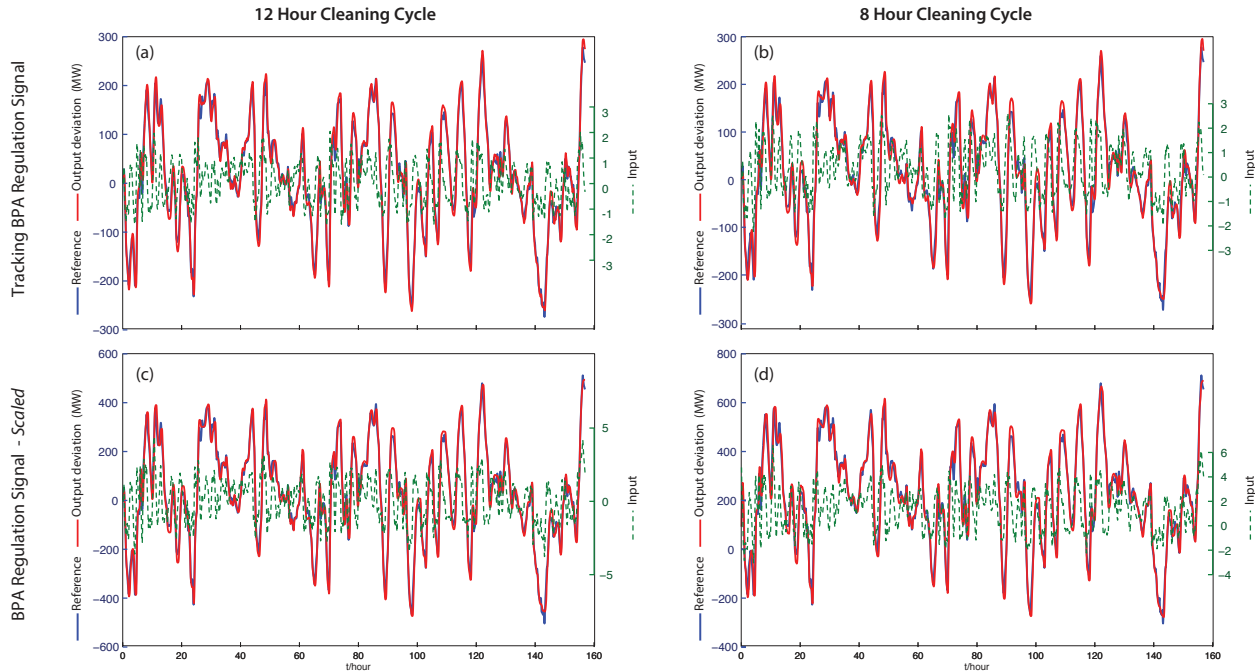


Fig. 11. Closed loop simulation in two scenarios, using two different reference signals.

Results from four experiments are shown in Fig. 11. Subplots (a) and (b) show tracking results using the low-pass filtered signal shown in Fig. 1, and the second row shows tracking performance when the signal magnitude is increased and shifted to match its potential capacity. The tracking performance is remarkable in all cases. In particular, it is surprising that a ± 400 MW signal can be tracked, given that the average power consumption of the pools is 500 MW in Scenario 1.

Subplots (a) and (b) in Fig. 12 shows what happens when the reference signal exceeds capacity. Two sources of error are evident in these plots. First, the power deviation saturates when all of the 10^6 pools are turned off, or all are turned on. Secondly, large tracking errors are observed immediately after saturation. This is a consequence of memory in the PI controller – what is known as *integrator windup*. To solve this problem, the BA should truncate the regulation signal so that it does not exceed the values $\{+Demand, -Supply\}$. Subplots (c) and (d) in Fig. 12 use the same regulation signal used in (a), (b), but truncated to meet these capacity constraints. Once again, the tracking is nearly perfect.

Individual risk: These simulation experiments have focused on the service to the grid, and the accuracy of the mean-field model approximations. The fidelity of approximation is remarkable.

The next question to ask is, what happens to an individual pool? Because of constraints on the regulation signal, it is found in simulations that the average cleaning time for each pool owner is close to the target

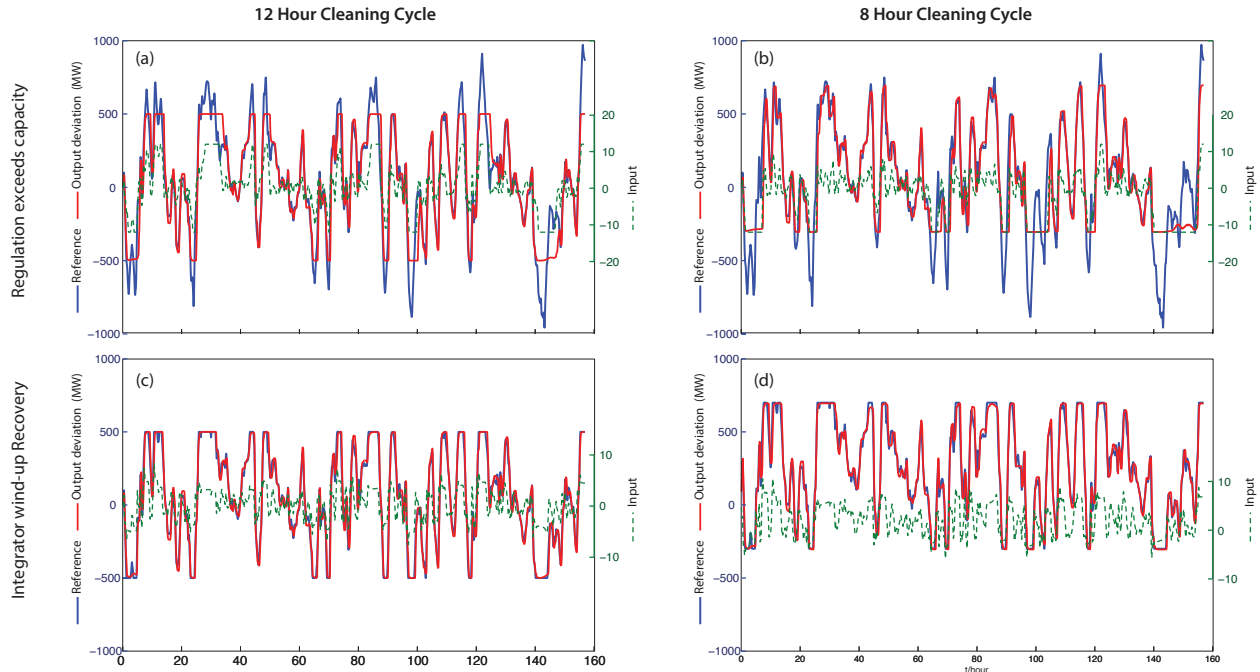


Fig. 12. The impact of exceeding capacity

values (either 12 or 8 hours per day in the two scenarios treated here). This is to be expected by the Law of Large Numbers.

The Central Limit Theorem can be appealed to if we wish to understand the impact of this control architecture on an individual pool. In simulations we find that the empirical distribution of hours cleaned over a four day period appears to be roughly Gaussian, and hence there is a portion of pools that are under-cleaned, and another portion that receive too many hours of cleaning.

Risk to individual consumers can be reduced or eliminated by using an additional layer of control at the loads. If over a period of several days the system detects over or under-cleaning, then the control system will ignore the signal sent by the BA. The aggregate impact of this modification represents a small amount of un-modeled dynamics. In preliminary experiments we have seen virtually no impact on tracking; only a small reduction in capacity. Analysis of individual risk is a topic of ongoing research.

IV. CONCLUSIONS

The simplicity of the MDP solution, and the remarkable accuracy of the LTI approximation for the mean-field model makes this approach appealing for this and many related applications.

There are several issues that have not been addressed here:

- (i) We do not fully understand the potential cost to consumers in terms of energy, or risk in terms of rare events in which the pool is under- or over-cleaned. It is likely that hard constraints on performance can be put in place without impacting the analysis.
- (ii) Does the grid operator need to know the real-time power consumption of the population of pools? Probably not. The BA is interested in regulating frequency, and this may be the only measurements needed for harnessing ancillary service from these loads. The grid frequency passed through a band-pass filter could serve as a surrogate for the measurement y_t assumed in this paper. It may be valuable to have *two* measurements at each load: The BA command, and local frequency measurements.
- (iii) How can we engage consumers? The formulation of contracts with customers requires a better understanding of the value of ancillary service, as well as consumer preferences.

REFERENCES

- [1] (2013) FPL on call saving program. [Online]. Available: www.fpl.com/residential/energy_saving/programs/oncall.shtml
- [2] J. MacDonald, P. Cappers, D. S. Callaway, and S. Kiliccote, “Demand response providing ancillary services a comparison of opportunities and challenges in the U.S. wholesale markets,” in *Grid-Interop*, Irving, TX, December 2012 2012.
- [3] (2013) BPA balancing authority load and total wind, hydro, and thermal generation (website). [Online]. Available: transmission.bpa.gov/business/operations/Wind/baltwg.aspx
- [4] H. Hao, A. Kowli, Y. Lin, P. Barooah, and S. Meyn, “Ancillary service for the grid via control of commercial building hvac systems,” in *American Control Conference*, June 2013.
- [5] 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*, 2012, pp. 1908–1913.
- [6] Y. Lin, P. Barooah, and S. Meyn, “Low frequency power grid ancillary service from commercial building HVAC systems,” in *Proc. IEEE Int. Conf. on Smart Grid Communications*, Oct. 2013, to appear.
- [7] S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp. 397–413, 1993.
- [8] R. Srikant, *The mathematics of Internet congestion control*, ser. Systems & Control: Foundations & Applications. Boston, MA: Birkhäuser Boston Inc., 2004.
- [9] E. Todorov, “Linearly-solvable Markov decision problems,” in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. Platt, and T. Hoffman, Eds. Cambridge, MA: MIT Press, 2007, pp. 1369–1376.
- [10] I. Kontoyiannis and S. P. Meyn, “Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes,” *Electron. J. Probab.*, vol. 10, no. 3, pp. 61–123 (electronic), 2005.
- [11] C. S. B. U. Design & Engineering Services. (2008, June) Pool pump demand response potential: Demand and run-time monitored data, DR 07.01 Report. [Online]. Available: www.etcc-ca.com/images/stories/pdf/ETCC_Report_473.pdf
- [12] K. Hall and M. Lo, “Survey of residential swimming pools assessed by Florida county property appraisers, Summer 2006,” Florida Department of Health, Tech. Rep., 2006.

- [13] J. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *IEEE Trans. on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.
- [14] Z. Ma, D. Callaway, and I. Hiskens, “Decentralized charging control for large populations of plug-in electric vehicles: Application of the Nash certainty equivalence principle,” in *Control Applications (CCA), 2010 IEEE International Conference on*, sept. 2010, pp. 191–195.
- [15] D.-C. Tomozei and J.-Y. L. Boudec, “Satisfiability of elastic demand in the smart grid,” *CoRR*, vol. abs/1011.5606, 2010.
- [16] R. Couillet, S. Perlaza, H. Tembine, and M. Debbah, “Electrical vehicles in the smart grid: A mean field game analysis,” *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 6, pp. 1086–1096, 2012.
- [17] M. Huang, P. E. Caines, and R. P. Malhame, “Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε -Nash equilibria,” *IEEE Trans. Automat. Control*, vol. 52, no. 9, pp. 1560–1571, 2007.
- [18] V. Borkar and R. Sundaesan, “Asymptotics of the invariant measure in mean field models with jumps,” *Stochastic Systems*, vol. 2, no. 2, pp. 322–380, 2012.
- [19] N. Gast, B. Gaujal, and J.-Y. Le Boudec, “Mean field for Markov decision processes: From discrete to continuous optimization,” *IEEE Trans. Automat. Control*, vol. 57, no. 9, pp. 2266–2280, 2012.
- [20] J. Unnikrishnan, D. Huang, S. P. Meyn, A. Surana, and V. V. Veeravalli, “Universal and composite hypothesis testing via mismatched divergence,” *IEEE Trans. Inform. Theory*, vol. 57, no. 3, pp. 1587–1603, 2011.
- [21] A. Dembo and O. Zeitouni, *Large Deviations Techniques And Applications*, 2nd ed. New York: Springer-Verlag, 1998.
- [22] 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 (Invited),” in *Proceedings of the 52nd IEEE Conf. on Decision and Control*, 2013.
- [23] S. P. Meyn and R. L. Tweedie, *Markov chains and stochastic stability*, 2nd ed. Cambridge: Cambridge University Press, 2009, published in the Cambridge Mathematical Library. 1993 edition online.
- [24] E. Seneta, *Non-Negative Matrices and Markov Chains*, 2nd ed. New York, NY: Springer, 1981.
- [25] E. Nummelin, *General Irreducible Markov Chains and Nonnegative Operators*. Cambridge: Cambridge University Press, 1984.
- [26] S. P. Meyn, *Control Techniques for Complex Networks*. Cambridge: Cambridge University Press, 2007, pre-publication edition available online.
- [27] V. S. Borkar, “Convex analytic methods in Markov decision processes,” in *Handbook of Markov decision processes*, ser. Internat. Ser. Oper. Res. Management Sci. Boston, MA: Kluwer Acad. Publ., 2002, vol. 40, pp. 347–375.
- [28] C. McDiarmid, “Concentration,” in *Probabilistic methods for algorithmic discrete mathematics*, ser. Algorithms Combin., M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, Eds. Berlin: Springer, 1998, vol. 16, pp. 195–248.
- [29] P. J. Schweitzer, “Perturbation theory and finite Markov chains,” *J. Appl. Prob.*, vol. 5, pp. 401–403, 1968.

APPENDIX

A. *Controlled transition matrix: Proof of Proposition 2.2.*

The dependency of η^* , $h = \log(v)$, and λ on ζ will be suppressed to simplify notation in the proof of Proposition 2.2 that follows.

The existence of a unique maximal eigenvalue and a positive eigenvector is a consequence of the Perron-Frobenius Theorem [24], [25]. Recent results on this and associated multiplicative ergodic theory are contained in [10], from which the identity $\Lambda := \log(\lambda) = \eta_\zeta^*$ can be established (see [10, Theorem 1.2]).

The bounds in (10) imply a rate of convergence of the finite horizon cost using \check{P} to its infinite-horizon limit, which in particular implies that $\Lambda := \log(\lambda) = \eta_\zeta^*$. To complete the proof of Proposition 2.2 it remains to establish this pair of bounds.

We begin with the second bound in (10). Let \check{p}^T denote the probability on strings on X^T induced by \check{P} , with initial condition x_0 given. This can be expressed,

$$\check{p}^T(x_1, \dots, x_T) = \frac{v(x_T)}{v(x_1)} \exp\left(\zeta \sum_{t=1}^T \mathcal{U}(x_t) - T\Lambda\right) p_0(x_1, \dots, x_T)$$

Since \check{p}^T is a probability measure we have,

$$1 = \sum_{x_i} \check{p}^T(x_1, \dots, x_T) = \sum_{x_i} \frac{v(x_T)}{v(x_1)} \exp\left(\zeta \sum_{t=1}^T \mathcal{U}(x_t) - T\Lambda\right) p_0(x_1, \dots, x_T)$$

which gives,

$$T\Lambda = \log\left\{\mathbb{E}\left[\frac{v(X_T)}{v(X_1)} \exp\left(\zeta \sum_{t=1}^T \mathcal{U}(X_t)\right)\right]\right\}.$$

The bound on $|T\Lambda - \mathcal{W}_T(p^{T*})|$ follows from this identity combined with Proposition 2.1 (which establishes $\mathcal{W}_T(p^{T*}) = \Lambda_T(\zeta)$).

Substitution of \check{p}^T into the definition of \mathcal{W}_T gives,

$$\mathcal{W}_T(\check{p}^T) = -\mathbb{E}_p\left[\log(v(X(T))) - \log(v(X(1)))\right] + T\Lambda$$

This combined with the second bound in (10) gives the desired upper bound on $\mathcal{W}_T(p^{T*}) - \mathcal{W}_T(\check{p}^T)$. \square

B. *Computation*

The approximations of η_ζ^* and h_ζ^* are based on Poisson's equation. In general, we say that g is the solution to Poisson's equation with forcing function f if the following equation holds:

$$g - P_0 g = f \tag{41}$$

For this finite state space Markov chain, it follows from invariance of π_0 that $\pi_0(f) = \sum_x \pi_0(x)f(x) = 0$.

We show here that the second order term \mathcal{S} defined in (17) can be expressed as a solution to a certain Poisson equation, and explain how to compute solutions.

1) *Poisson's equation:* This is a finite state Markov chain, so Perron-Frobenius theory is very simple and attractive [10], [24], [25]. Let $s: \mathsf{X} \rightarrow \mathbb{R}_+$ be a function (not identically zero), and ν a probability measure such that the *minorization condition* holds,

$$P_0(x, y) \geq s(x)\nu(y), \quad x, y \in \mathsf{X}$$

This is written $P \geq s \otimes \nu$. We will take $s(x) = \mathbb{I}\{x = \alpha\}$ and $\nu(y) = P_0(\alpha, y)$, so that for all y ,

$$\bar{P}(x, y) := P_0(x, y) - s(x)\nu(y) = \begin{cases} 0 & x = \alpha \\ P_0(x, y) & x \neq \alpha \end{cases}$$

The invariant probability measure π_0 can be expressed in terms of a matrix inverse. We first use the implication,

$$\pi_0 P_0 = \pi_0 \implies \pi_0 (P_0 - s \otimes \nu) = \pi_0 - \delta \mu$$

where $\delta = \sum_x \pi_0(x)s(x)$ is a constant. We then invert,

$$\pi_0 = \delta \nu [I - (P_0 - s \otimes \nu)]^{-1}$$

We substitute $\bar{P} = P_0 - s \otimes \nu$, and note that the measure $\mu = \nu [I - \bar{P}]^{-1}$ is the unnormalized invariant measure. Normalization gives the invariant probability measure,

$$\pi_0(x) = \mu(x) / [\sum_y \mu(y)], \quad x \in \mathsf{X}.$$

The inverse $Z = [I - \bar{P}]^{-1}$ is called the potential matrix [25].

A similar computation is used to solve Poisson's equation. Assume that $\nu(H) = \sum \nu(x)H(x) = 0$. This is without loss of generality. Under this assumption,

$$\tilde{U} + P_0 H = H \implies \tilde{U} + (P_0 - s \otimes \nu) H = H$$

which gives $H = Z\tilde{U}$.

2) *Second-order approximation:* Let h'_ζ and h''_ζ denote the first and second derivatives of h_ζ^* with respect to the argument ζ . Each of these are real-valued functions on X . The Taylor series approximation (21) in Proposition 2.3 implies that

$$h'_\zeta \Big|_{\zeta=0} = H \quad \text{and} \quad h''_\zeta \Big|_{\zeta=0} = \mathcal{S}.$$

Here we obtain representations of these functions that are convenient for computation.

We show in Proposition A.1 that \mathcal{S} solves a certain Poisson equation, that can be solved through elementary matrix algebra. The representation for h''_ζ in Proposition A.1 appears to be new.

The derivations are most easily obtained using the *nonlinear generator*, defined for any function $g: \mathsf{X} \rightarrow \mathbb{R}$ via,

$$\mathcal{H}(g) := \mathcal{P}(g) - g,$$

where $\mathcal{P}(g)$ denotes the function,

$$\mathcal{P}(g)\Big|_x := \log(Pe^g)\Big|_x = \log\left(\sum_y P(x, y)e^{g(y)}\right), \quad x \in \mathsf{X}$$

Using this notation, the eigenvector equation can be expressed in a form similar to Poisson's equation,

$$h_\zeta^* - \mathcal{P}(h_\zeta^*) = z\mathcal{U} - \eta_\zeta^* \quad (42)$$

in which $z\mathcal{U} - \eta_\zeta^*$ plays the role of a forcing function, similar to (41). The similarity between (41) and (42) is why h_ζ^* is called the solution to the *multiplicative Poisson equation* in [10].

We also require a nonlinear operator that defines one-step variance: For any function $g: \mathsf{X} \rightarrow \mathbb{R}$,

$$\mathcal{V}(g)\Big|_x := \sum_y P_0(x, y)g^2(y) - \left[\sum_y P_0(x, y)g(y)\right]^2$$

That is, $\mathcal{V}(g) = Pg^2 - [Pg]^2$.

Proposition A.1: The functions h'_ζ and h''_ζ solve the following Poisson equations:

$$\mathcal{U} + \check{P}_\zeta h'_\zeta = h'_\zeta + \Lambda'_\zeta \quad (43)$$

$$\mathcal{V}(h'_\zeta) + \check{P}_\zeta h''_\zeta = h''_\zeta + \Lambda''_\zeta \quad (44)$$

where $\Lambda_\zeta = \log(\lambda_\zeta) = \eta_\zeta^*$. The derivatives also satisfy the boundary conditions,

$$h'_\zeta(\boldsymbol{\alpha}) = h''_\zeta(\boldsymbol{\alpha}) = 0.$$

In particular, with $\zeta = 0$ we obtain,

$$\mathcal{U} + P_0 H = H + \eta_0 \quad (45)$$

$$\mathcal{V}(H) + P_0 \mathcal{S} = \mathcal{S} + \kappa^2 \quad (46)$$

□

Proof: The boundary conditions hold because $h_\zeta^*(\boldsymbol{\alpha}) = \log(v_\zeta(\boldsymbol{\alpha})) = 0$ for all ζ . This is by construction — see (13) and (12).

The remaining results are obtained on differentiating each side of (42). The first derivative of $\mathcal{P}(h_\zeta^*)$ is obtained from simple calculus:

$$\frac{d}{d\zeta}\mathcal{P}(h_\zeta^*) = \frac{d}{d\zeta}\log(Pe^{h_\zeta^*}) = \frac{1}{Pe^{h_\zeta^*}}P\{e^{h_\zeta^*}h'_\zeta\} = \check{P}_\zeta h'_\zeta \quad (47)$$

On differentiating each side of (42) we thus obtain (43).

For the second derivative of h_ζ^* we require the second derivative of $\mathcal{P}(h_\zeta^*)$. A representation will follow from the product rule, using (47):

$$\frac{d^2}{d\zeta^2}\mathcal{P}(h_\zeta^*) = \check{P}_\zeta h''_\zeta + \check{P}'_\zeta h'_\zeta \quad (48)$$

To obtain an expression for the matrix \check{P}'_ζ , observe that for any $x, y \in \mathsf{X}$ for which $P_0(x, y) > 0$,

$$\begin{aligned} \frac{1}{\check{P}_\zeta(x, y)} \frac{d}{d\zeta} \check{P}_\zeta(x, y) &= \frac{d}{d\zeta} \log(\check{P}_\zeta(x, y)) \\ &= \frac{d}{d\zeta} \log(P(x, y)e^{h_\zeta^*(y)} / Pe^{h_\zeta^*(x)}) \\ &= h'_\zeta(y) - \frac{d}{d\zeta} \mathcal{P}(h_\zeta^*) \Big|_x = h'_\zeta(y) - \check{P}_\zeta h'_\zeta(x) \end{aligned}$$

where the final equality follows from (47). This gives,

$$\frac{d}{d\zeta} \check{P}_\zeta(x, y) = \check{P}_\zeta(x, y)[h'_\zeta(y) - \check{P}_\zeta h'_\zeta(x)], \quad x, y \in \mathsf{X}.$$

Substituting this into (48) gives,

$$\frac{d^2}{d\zeta^2} \mathcal{P}(h_\zeta^*) \Big|_x = \check{P}_\zeta h''_\zeta(x) + \check{P}'_\zeta h'_\zeta(x) = \check{P}_\zeta h''_\zeta(x) + \sum_y \check{P}_\zeta(x, y)[h'_\zeta(y) - \check{P}_\zeta h'_\zeta(x)]h'_\zeta(y)$$

which simplifies to,

$$\frac{d^2}{d\zeta^2} \mathcal{P}(h_\zeta^*) = \check{P}_\zeta h''_\zeta + \mathcal{V}(h'_\zeta) \quad (49)$$

Differentiating each side of (42) twice thus gives (44). \square