Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in Florida

Sean Meyn, Prabir Barooah, Ana Bušić, 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. We focus on pool pumps, and how they can be used to provide ancillary service to the grid for maintaining demand-supply balance. A Markovian Decision Process (MDP) model is introduced for an individual pool pump. 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 pools 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. Simulations are provided to illustrate the accuracy of the approximations and effectiveness of the proposed control approach.

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.

The research described in this paper concerns harnessing the flexibility of aggregates of deferrable 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.

In this paper we 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 [7]. 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 [8]. 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.

The technical content of the paper starts with a control architecture designed to address privacy concerns and communication constraints. An individual pool 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 pool. This avoids synchronization, much like randomized congestion avoidance protocols in communication networks.

To formulate a randomized control strategy, a Markovian Decision Process (MDP) model is proposed for an individual pool pump. An aggregate model for a large number of pools 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 fraction of pool pumps in the on state as the output. The output is therefore a measure of the aggregate demand. The LTI approximation is convenient for control design at the grid level: the scalar input becomes the control signal that the BA will broadcast to all the pools, which adjusts a parameter in the randomized policy for the optimal MDP solution at each pool. In the numerical results in this paper it is found that a PI (proportional-integral) controller works very well.

The analysis of the aggregate in Section II can be viewed as a mean-field game, as in the prior work [9], [10]. The LTI-approximation is based on large-deviation theory, made possible by applying a particular optimal control approach of Todorov [11] along with results from [12].

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 [9], [10], [13], [14] 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].

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 pool: An MDP control solution is obtained at each pool pump, 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 pools 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 the pool and the grid

In this section we pose an optimal control formulation for an individual pool, taking into account the needs of the pool and the grid. Approximations of the control solution are obtained and illustrated in examples. In Section III we analyze the aggregate behavior of a collection of pools each governed by the optimal policy.

A. Pool-Grid model

Fig. 2. State transition diagram for the pool-pump model.

The state process of the MDP model for an individual pool is denoted $\boldsymbol{X}$, which evolves on the state space

$$\mathsf{X} = \{(m,i) : m \in \{\oplus, \ominus\}, \ i \in \{0, \dots, T\}\};$$

$X(t) = (\ominus, i)$ 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. 2.

A randomized control solution is desirable to avoid synchronization of pools in the grid, and also to facilitate analysis of the aggregate system. This is achieved by imposing the following cost structure on the MDP model.

It is assumed that there is a transition matrix P_0 that models "control free" behavior of the MDP. A new transition matrix P may be chosen based on the state of the grid, but there is a cost for deviation if $P \neq P_0$. Moreover, there is a cost function κ on X that is used to model the needs of the grid. The one-stage cost for the MDP model is given by the weighted sum,

$$c(x, P) = z\kappa(x) + D(P(x, \cdot)\|P_0(x, \cdot)) \qquad (1)$$

where z is a scalar constant (for now), and where D denotes Kullback-Leibler divergence

$$D(P(x, \cdot)\|P_0(x, \cdot)) = \sum_y P(x, y) \log\Big(\frac{P(x, y)}{P_0(x, y)}\Big)$$

The cost will be infinite for any x and P for which $P(x, \cdot)$ is not absolutely continuous with respect to $P_0(x, \cdot)$.

The average-cost optimality criterion is considered here. The optimal average cost is defined as follows:

$$\eta_z^* := \min_{\pi, P}\Big\{\sum \pi(x)c(x, P) : \pi \text{ is invariant for } P\Big\} \quad (2)$$

The dependency of η^* on z will be suppressed to simplify notation. It was shown in [11] that the solution to the average cost optimality equation associated with (2) can be found through the solution of an eigenvector problem

$$\widehat{P}_0 v = \lambda v \qquad (3)$$

with $\lambda = e^{-\eta^*}$, and $\widehat{P}_0$ the scaled transition matrix,

$$\widehat{P}_0(x, y) = \exp(-z\kappa(x))P_0(x, y), \qquad x, y \in \mathsf{X}. \quad (4)$$

The optimizing transition matrix is

$$\check{P}(x, y) = \frac{1}{\lambda}\frac{1}{v(x)}\widehat{P}_0(x, y)v(y), \qquad x, y \in \mathsf{X}. \quad (5)$$

Second order Taylor series approximations for v and η^* near $z \sim 0$ can be found by borrowing tools from large-deviations theory. These approximation results are new, and are collected together in the Appendix.

We now fix the structure for $c(x, P)$ that is used in the numerical results described in this paper. 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. 2. In the numerical results described below a symmetric model was chosen for P_0 in which $p_i^{\oplus} = p_i^{\ominus}$, where $p_i^{\oplus} := P(\text{pump switches on} \mid \text{it has been off } i \text{ hours})$. Similarly, $p_i^{\ominus} := P(\text{pump switches off} \mid \text{it has been on } i \text{ hours})$. Fig. 3 shows a plot of the probability $p_i^{\oplus}$ vs. i.

Symmetry is *not* required in any of the theoretical results of this paper, such as the form of the solution to the optimality equations.

The cost function κ on X is chosen as the indicator function that the pool pump is not operating:

$$\kappa(x) = \sum_i \mathbb{I}\{x = (\ominus, i)\} \qquad (6)$$

The parameter z in (1) can be positive or negative; If $z > 0$ this one-stage cost function provides incentive to turn pumps on.


Fig. 3. Control free behavior of a pool used for numerical studies.

Given this structure, the solution to the eigenvector equation (3) has a simple form described in the Appendix (see (20) and (31)).

Approximations: A Taylor-series approximation of the average cost is based on two parameters, defined with respect to the nominal model P_0 with invariant measure π_0. The first-order coefficient is the the steady-state mean of κ,

$$\eta_0 = \sum_x \pi_0(x)\kappa(x) \qquad (7)$$

By symmetry, in this model we have $\eta_0 = 1/2$. The second-order coefficient is based on the *asymptotic variance* of κ for the nominal model (the variance appearing in the Central Limit Theorem for the nominal model). For this finite state space model this is represented as the limit,

$$\varsigma_0^2 = \lim_{N\to\infty}\frac{1}{N}\mathsf{E}\Big[\Big(\sum_0^{N-1}\tilde{\kappa}(X(t))\Big)^2\Big] \qquad (8)$$

where $\tilde{\kappa} = \kappa - \eta_0$.

The approximation of the average cost established in Proposition 1.1 in the Appendix is then,

$$\eta_z^* = \eta_0 z - \tfrac{1}{2}\varsigma_0^2 z^2 + O(z^3) \qquad (9)$$

Shown in Fig. 4 is a comparison of η_z^* with the quadratic approximation based on (9).

Fig. 4. The optimal cost η_z^* and its quadratic approximation.

The invariant measure (or steady-state distribution) for $\check{P}_z$ is denoted $\check{\pi}_z$. The steady-state probability that a pool-pump is in operation is given by

$$\check{P}\{\text{pool-pump is on}\} = 1 - \sum_x \check{\pi}_z(x)\kappa(x)$$

A linear approximation is obtained in Proposition 1.1 (ii):

$$\check{P}\{\text{pool-pump is on}\} = 1 - (\eta_0 - \varsigma_0^2 z) + O(z^2) \quad (10)$$

A comparison of the true probability and its affine approximation is shown in Fig. 5. The approximation is very tight for $|z| \leq 3$. For larger values of z the true steady-state probability saturates (approximately 0.9 as $z \to +\infty$).

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

For fixed z, 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 $z = 0, \pm 4, \pm 6$. The plots of $\check{p}^{\ominus}$ are obtained through symmetry.


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

III. CONTROLLING A LARGE NUMBER OF POOLS

Consider N pool-pumps operating independently under the randomized policy described in the previous section. The state of the ith pool is denoted $X^i(t)$. For large N we have from the Law of Large Numbers,

$$\frac{1}{N} \sum_{i=1}^{N} \kappa(X^i(t)) \approx \mathsf{E}[\kappa(X(t))] \quad (11)$$
$$= \mathsf{P}\{\text{Pump not operating at time } t\}$$

The expectation and probability on the right are with respect to the optimal transition law $\check{P}_z$, where z is the cost parameter used in (1).

There is now a centralized control problem faced by the BA: How to choose the variable z to regulate the power consumed by the pool pumps?

A. Mean-field control model

The BA observes the proportion of pools that are on as a function of time. Based on this information and the state of the grid, the command z is computed at the BA, and transmitted to each pool pump. The control at each pool pump is decentralized, based only on its own state and the signal z.

To address the control problem faced by the BA it is necessary to relax the assumption that this parameter is fixed. We let $z = \{z_0, z_1, \dots\}$ 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 it is assumed without justification the existence of a mean-field limit: Let $N \to \infty$ to obtain the generalization of (11),

$$\lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}\{X^i(t) = x\} = \mu_t(x), \quad x \in \mathsf{X}. \quad (12)$$

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

$$\sum_{x_i \in \mathsf{X}} \mu_0(x_0) \check{P}_{z_0}(x_0, x_1) \check{P}_{z_1}(x_1, x_2) \cdots \check{P}_{z_{t-1}}(x_{t-1}, x_t)$$

where x_t is an arbitrary state in X, and the sum is over all intermediate states. We view $\{\mu_t\}$ as a state process that is under our control through z. The grid operator is interested in the average number of pools pumps that are operating, given by

$$y_t = 1 - \sum_x \mu_t(x) \kappa(x)$$

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. Further details are given in Section III-C.

In the linear model we let Φ_t denote a column vector of dimension $|\mathsf{X}|$ whose xth component approximates $\mu_t(x)$, and γ_t a scalar that approximates y_t. The approximate dynamics of $\{\mu_t, y_t\}$ are defined by a linear state space model,

$$\Phi_{t+1} = A\Phi_t + Bz_t$$
$$\gamma_t = C\Phi_t \quad (13)$$

with initial condition $\Phi_0(x) = \mu_0(x)$.

We begin with an equilibrium analysis in which z is held constant.

B. Equilibrium

Suppose that z does not vary with time, $z_t = z^*$ for all t, and consider the steady-state behavior of the mean-field model. We denote $y_\infty = \lim_{t \to \infty} y_t$, which is the steady-state probability that a pool is on, for the model with transition law $\check{P}_{z^*}$. This can be approximated using Proposition 1.1:

$$y_\infty = \mathsf{P}\{\text{Pump is operating}\} \approx 1 - (\eta_0 - \varsigma_0^2 z^*)$$

From the viewpoint of the BA, there is a value G^* of desired consumption by all the pools. If $g_\mathrm{p} > 0$ denotes the

Fig. 7. Frequency response for linearized model $C[Iz - A]^{-1}B$

consumption of one pool pump in operation, and if there are N pools in total, then the desired steady-state probability is $y_\infty = G^*/(Ng_\mathrm{p})$. This translates to a corresponding value of z^*,

$$z^* \approx \frac{1}{\varsigma_0^2}\left[\frac{1}{g_\mathrm{p}}\frac{G^*}{N} - (1 - \eta_0)\right] = \frac{1}{\varsigma_0^2}\frac{1}{g_\mathrm{p}}\frac{\widetilde{G}}{N} \qquad (14)$$

where $\widetilde{G} = G^* - G_0$, with $G_0 = g_\mathrm{p}N(1 - \eta_0)$, the control-free value obtained with $z^* = 0$.

C. Linear systems analysis

Consider now the case in which z is a function of time. We construct a linear systems approximation in which this is the input, and the output is y (the fraction of pools on).

In the state space model (13) we take $A = P_0^\tau$, and C is the row vector of dimension $|\mathsf{X}|$ defined so that $\sum_x C_x\mu_t(x) = 1 - \sum_x \mu_t(x)\kappa(x) = y_t$. Specification of B is obtained next, based on the Taylor series approximations developed in Section B of the Appendix.

We have, from (5),

$$\check{P}_z(x, y) = e^{-z\kappa(x)+\eta_z^* - \mathcal{V}_z(x)}P_0(x, y)e^{\mathcal{V}_z(y)}$$

Based on the first order approximation of $\mathcal{V}_z$ in Proposition 1.1 we obtain,

$$\check{P}_z(x, y) \approx e^{z[h_0(x)-\tilde{\kappa}(x)]}P_0(x, y)e^{-zh_0(y)}$$

where h_0 is a solution to *Poisson's equation* with forcing function κ, for the nominal model (see (21)). Using a first order Taylor series for the exponential then gives,

$$\check{P}_z(x, y) \approx [1 + z(h_0(x) - \tilde{\kappa}(x))]P_0(x, y)[1 - zh_0(y)]$$
$$\approx P_0(x, y) + z\mathcal{E}(x, y)$$

where $\mathcal{E}(x, y) = \big[(h_0(x) - \tilde{\kappa}(x))P_0(x, y) - P_0(x, y)h_0(y)\big]$.

If $\mu \sim \pi_0$ (the invariant measure for P_0) and z is small, then we can approximate,

$$\mu\check{P}_z \approx \mu P_0 + zB^\tau$$

where B is the column vector with entries $B(y) = \sum_x \pi_0(x)\mathcal{E}(x, y)$. In this way we can approximate y as the output γ of the linear system (13).

Observe that the linear system (13) is not the form of a typical linearization in which the state represents deviation from nominal behavior (in this case this is $z \equiv 0$, $\Phi := \pi_0^\tau$). We obtain a linearized system with $\Phi_t(x) \approx \mu_t(x)$, rather

than the deviation $\mu_t(x) - \pi_0(x)$, because the matrix A has a unique eigenvalue at 1, with eigenvector π_0^τ.

Fig. 7 shows the Bode plot for the linear model (13), and its pole-zero plot is shown in Fig. 8. The transfer function from z to γ is BIBO stable and minimum phase.

Fig. 9 shows remarkable solidarity between the Markov model and its linearization. The input was chosen to be swept-sinuosoid, of the form

Fig. 8. Open-loop pole-zero plot

$$z_t = g\sin(\delta t + \varepsilon t^2),$$

with g, δ, ε positive and small numbers.

D. Control design and simulation

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 dispatachable 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. Two mechanisms are used together to improve performance.

First, in the experiments conducted it was found that sampling at 30 minute intervals provided better response to the grid, with little impact on service to the network of pools.

The second mechanism is more subtle. Why should each pool operate on the same clock? We assume that each pool checks the regulation signal each $T = 30$ minutes. However, the pools have no common clock. It is convenient to model this via binning of time. For example, 1/6 of the pools measure the regulation signal at the top of the hour, 1/6 measure at five minutes past, and so on.

To model this refinement, let $m > 1$ denote a "super-sampling" parameter, and we fix $T = 30$ minutes for the sampling time of each pool. The model is in discrete time, with sampling interval T/m. For example, $m = 6$ 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 ith class that are operating, the total that are operating at time t is the sum $y_t = \sum_i y_t^i$. The delay in this model is T/m rather than T.

Fig. 9. Response to input $z_t = 0.3 \sin(\delta t + \varepsilon t^2)$ for Markov model and its linear approximation. There is a strong resonance at a frequency with period approximately 24 hours.

A linear model is minimum phase as before, but now only marginally since the sum introduces zeros on the unit circle. Regardless, it is found that high-gain control is suitable for the linear model.

For the purposes of translation to megawatts, it is assumed that there are $N = 10^6$ pools (consistent with the state of Florida), and that each pool in operation consumes $g_p = 1$ KW. Power consumption at time t is assumed to be equal to $Ng_p y_t$ (in KW), or $10^3 y_t$ (in MW).

The following experiments involve the mean-field Markovian model without linearization using PI control.[1] With a value of $m = 12$, corresponding to 150 second sampling intervals at the grid level, a proportional gain of 20, and integral gain of 4 worked well.


Fig. 10. Closed loop simulation with Markovian model.

In the results shown in Fig. 10 the reference signal was chosen to be the BPA regulation signal passed through a low pass filter. The tracking performance is remarkable.

Observe that the magnitude of service reaches 200 MW. This is a substantial amount of ancillary service in this frequency band, and there is no reason to believe that a network of one million pools cannot supply much more.

IV. Conclusions

The simplicity of the MDP solution, and the remarkable accuracy of the LTI approximation for the mean-field model

---

[1]Just prior to final submission, experiments using a simulation of $100,000$ pools were completed by Yue Chen, a graduate student at Univ. of Florida. The results are nearly identical to those described here, using the deterministic mean-field game model, and will appear on our website and a journal article in preparation.

makes this approach appealing for this and many related applications.

There are also many issues that have not been addressed here: The formulation of contracts with customers requires a better understanding of the value of ancillary service, as well as consumer preferences. Moreover, we do not yet 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 – this is a focus of current research.

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 us 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] 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

[8] 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.

[9] R. Couillet, S. Perlaza, H. Tembine, and M. Debbah, "A mean field game analysis of electric vehicles in the smart grid," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*, march 2012, pp. 79 –84.

[10] 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.

[11] 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.

[12] 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.

[13] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *Power Systems, IEEE Transactions on*, vol. 28, no. 1, pp. 430 –440, feb. 2013.

[14] D.-C. Tomozei and J.-Y. L. Boudec, "Satisfiability of elastic demand in the smart grid," *CoRR*, vol. abs/1011.5606, 2010.

[15] 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.

[16] S. P. Meyn, *Control Techniques for Complex Networks.* Cambridge: Cambridge University Press, 2007, pre-publication edition available online.

[17] A. Dembo and O. Zeitouni, *Large Deviations Techniques And Applications*, 2nd ed. New York: Springer-Verlag, 1998.

[18] 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.

[19] E. Seneta, *Non-Negative Matrices and Markov Chains*, 2nd ed. New York, NY: Springer, 1981.

[20] E. Nummelin, *General Irreducible Markov Chains and Nonnegative Operators.* Cambridge: Cambridge University Press, 1984.

[21] 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.

APPENDIX

AVERAGE COST OPTIMAL CONTROL: AN INFORMATION-THEORETIC SETTING

This Appendix consists of three parts: A short review of Todorov's eigenvector characterization for a class of MDPs, some new Taylor-series approximations for the solution of this optimal control problem, and specialized results for the pool-pump MDP model introduced in this paper. The control setting has an information-theoretic flavor since the cost is based on deviation from nominal behavior, where the deviation-cost is based on relative entropy. The analysis here is similar to what is used in information theory and statistics (in particular [12], [15]).

A. Todorov's formulation

We first review a general control methodology introduced by Todorov in [11]. We present here the main ideas; modifying the original formulation to simplify discussion.

In the standard formulation of optimal control of Markov chains – also called Markov decision theory – there is a finite state space X and action space U. A state process $\boldsymbol{X}$ evolves on X, and input process $\boldsymbol{U}$ evolves on U. A controlled Markov transition matrix $P_u(x,y)$ defines the probability that the state moves from x to y in one step:

$$\mathsf{P}\{X(t+1) = y \mid X(t) = x, U(t) = u\} = P_u(x,y). \quad (15)$$

Under general conditions [16], the optimal average cost η^* is independent of the initial condition, and together with a *relative value function* h^* solves the *average-cost dynamic programming equation* (ACOE),

$$\min_u\{c(x,u) + P_u h^*(x)\} = h^*(x) + \eta^* \quad (16)$$

In this equation we use the standard shorthand,

$$P_u h^*(x) = \sum_y P_u(x,y)h^*(y)$$
$$= \mathsf{E}[h^*(X(t+1)) \mid X(t) = x, U(t) = u]$$

There are two components to the formulation of [11]. First, the action space U is identified with the set of all transition matrices (restrictions are placed on permissible P_u in [11], but this is inconsequential here). This is consistent with the definition (15), in which on observing $X(t) = x$, the entire probability law $P(x,y)$ is determined at this time. Second, the cost function is taken to be of the specific form (1) for some function $\kappa \colon \mathsf{X} \to \mathbb{R}$, a scalar z, and nominal transition matrix P_0.

The ACOE (16) becomes a minimum over P,

$$\min_P\{c(x,P) + Ph^*(x)\} = h^*(x) + \eta^* \quad (17)$$

This convex optimization problem appears in the theory of large deviations of sample-means [17]. In the minimization (17) we can fix x, and let $\mu_0(\cdot) = P_0(x, \cdot)$, so that the optimization problem is over probability vectors $\mu = P(x, \cdot)$ on X. Given the form of the cost function, the minimization to be solved is,

$$\min_\mu\{D(\mu\|\mu_0) + \mu(h^*)\} = -\max_\mu\{\mu(-h^*) - D(\mu\|\mu_0)\}$$

where the maximum is of the standard form seen in large deviations for i.i.d. stochastic processes. The solution $\check{\mu}$ is the "twisted distribution",

$$\check{\mu}(y) = \xi\mu_0(y)e^{-h^*(y)}, \qquad y \in \mathsf{X}.$$

where $\xi > 0$ is a normalizing constant, chosen so that $\sum_y \check{\mu}(y) = 1$.

Returning to the original notation, and recalling that we solve this optimization problem for each x, the optimizer $\check{P}$ has the form,

$$\check{P}(x,y) = \xi(x)P_0(x,y)e^{-h^*(y)} \quad (18)$$

where the normalizing constant $\xi(x)$ is found as follows,

$$\frac{1}{\xi(x)} = \sum_y P_0(x,y)e^{-h^*(y)}.$$

Substituting this expression into (17) we find that the solution to the Perron-Frobenious eigenvector problem (3) is obtained with $v = e^{-h^*}$ and $\lambda = e^{-\eta^*}$.

Under the assumption that P_0 is irreducible, λ is the unique maximal eigenvalue for $\widehat{P}_0$, and v is unique up to constant multiples. This is not surprising, since the relative value function h^* solving the ACOE is typically unique up to an additive constant.

B. Representations and approximations

A solution to the eigenvector problem (3) can be represented through a regenerative formula. Let $\boldsymbol{\alpha} \in \mathsf{X}$ 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 $\boldsymbol{\alpha}$-irreducible [18]. 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,

$$\tau = \min\{t \geq 1 : X(t) = \boldsymbol{\alpha}\}.$$

Recall that the optimal cost, which of course depends upon z, is given by $\eta_z^* = -\log(\lambda)$. From the theory of positive matrices [12], [19], [20], it follows that it is the unique solution to,

$$1 = \mathsf{E}_\alpha\left[\exp\left(\sum_0^{\tau-1}[-z\kappa(X(t)) + \eta_z^*]\right)\right] \quad (19)$$

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

$$v(x) = \mathsf{E}_x\left[\exp\left(\sum_0^{\tau-1}[-z\kappa(X(t)) + \eta_z^*]\right)\right] \quad (20)$$

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

Based on these two formula and Taylor series expansions of the exponential, we can obtain second-order Taylor series approximations of η_z^* and $\mathcal{V}_z = \log(v_z)$ for $z \sim 0$.

Both approximations require the solution to *Poisson's equation* for P_0,

$$h_0(x) = \mathsf{E}_x\left[\sum_0^{\tau-1}\tilde\kappa(X(t))\right] \quad (21)$$

where $\tilde\kappa = \kappa - \eta_0$, and η_0 is the nominal steady-state mean (7). The asymptotic variance associated with κ defined in (8) can be expressed in terms of Poisson's equation [16], [18]:

$$\varsigma_0^2 = \sum_x \pi_0(x)\big(\tilde\kappa(x)h_0(x) - \tilde\kappa(x)^2\big)$$

The second order approximation of the optimal cost is given in terms of these two statistics. The approximation of the relative value function h^* is based on another variance term,

$$\mathcal{S}_0(x) = \mathsf{E}_x\left[\left(\sum_0^{\tau-1}\tilde\kappa(X(t))\right)^2\right] - \big(h_0(x)\big)^2 \quad (22)$$

Proposition 1.1: The following hold for the finite state space model in which P_0 is irreducible:

(i) The optimal cost η_z^* is concave as a function of z, and admits the Taylor series expansion,

$$\eta_z^* = \eta_0 z - \tfrac{1}{2}\varsigma_0^2 z^2 + O(z^3) \quad (23)$$

(ii) The mean of κ under the the invariant measure $\check\pi_z$ for $\check P_z$ is given by,

$$\sum_x \check\pi_z(x)\kappa(x) = \frac{d}{dz}\eta_z^* \quad (24)$$

This admits the first-order Taylor series approximation

$$\frac{d}{dz}\eta_z^* = \eta_0 - \varsigma_0^2 z + O(z^2) \quad (25)$$

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

$$h^*(x) = zh_0(x) - \tfrac{1}{2}z^2\mathcal{S}_0(x) + O(z^3) \quad (26)$$

Proof: Equations (23)—(25) follow from the fact that $-\eta_z^* = \log(\lambda)$ can be expressed as a cumulative log-moment generating function [12, Prop. 4.9].

Concavity follows from the fact that η_z^* is the minimum of linear functions of z (following the linear-program formulation of the ACOE [21]).

Next we establish the approximation of $\mathcal{V} = \log(v)$. The second order approximation of $\mathcal{V}(x)$ follows from the representation (20),

$$\mathcal{V}(x) = -zh_0(x) + \tfrac{1}{2}z^2\mathcal{S}_0(x) + O(z^3) \quad (27)$$

This gives an approximation for the relative value function (26) using $v = e^{-h^*}$. ∎

C. Solution to the pool pump MDP control problem

The representation of the eigenvector v given in (20) is based on the first return time τ to a state $\boldsymbol{\alpha} \in \mathsf{X}$. A convenient choice in the pool-pump model is $\boldsymbol{\alpha} = (\oplus, 0)$, so that τ is the first time that the pool-pump is switched from off to on,

$$\tau = \tau_\oplus = \min\{t \geq 1 : X(t) = (\oplus, 0)\}$$

The function v given in (20) can be found via dynamic programming.

Consider first an initial condition in which the pool pump is initially *off*. For $x = X(0) = (\ominus, i)$ we then have

$$v(x) = \mathsf{E}_x\big[\exp(\theta\tau_\oplus)\big], \quad (28)$$

with $\theta = \eta_z^* - z$. The case $i = T$ has an explicit solution, which is interpreted as a boundary condition,

$$v(x) = e^\theta, \qquad x = (\ominus, T). \quad (29)$$

This holds because T is an upper limit on the length of a cycle, so that $p_T^\oplus = 1$.

For $x = X(0) = (\ominus, i)$ with $i < T$, the next state can take on two possible values: $X(1) = (\ominus, i+1)$ or $X(1) = (\oplus, 0)$. In the latter case $\tau_\oplus = 1$, giving the recursive equation,

$$v(\ominus, i) = p_i^\oplus e^\theta + (1 - p_i^\oplus)e^\theta v(\ominus, i+1) \quad (30)$$

Consider now the initial condition $x = (\oplus, i)$ for $i \geq 0$. To reach the state $(\oplus, 0)$ for $t \geq 1$ it is necessary to first pass through the state $(\ominus, 0)$. On denoting

$$\tau_\ominus = \min\{t \geq 1 : X(t) = (\ominus, 0)\}$$

the value function can be represented,

$$v(x) = \mathsf{E}_x[e^{\eta^*\tau_\ominus}]\mathsf{E}_{(\ominus,0)}[e^{\theta\tau_\oplus}] \quad (31)$$

This is because $\kappa(X(t)) = 0$ for $t < \tau_\ominus$.

The moment generating function $\mathsf{E}_x[e^{\eta^*\tau_\ominus}]$ is easily computed as before.

Recall that Proposition 1.1 contains quadratic approximations for η_z^* and $\mathcal{V}_z = \log(v_z)$, based on the solution to Poisson's equation. This expectation defining h_0 in (21) is is similar to the representation of v, so that similar dynamic programming steps can be used to compute h_0.