

Individual Risk in Mean Field Control with Application to Automated Demand Response

Yue Chen, Ana Bušić, and Sean Meyn

Abstract—Flexibility of energy consumption can be harnessed for the purposes of ancillary services in a large power grid.

In prior work by the authors a randomized control architecture is introduced for individual loads for this purpose. In examples it is shown that the control architecture can be designed so that control of the loads is easy at the grid level: Tracking of a balancing authority reference signal is possible, while ensuring that the quality of service (QoS) for each load is acceptable *on average*. The analysis was based on a mean field limit (as the number of loads approaches infinity), combined with an LTI-system approximation of the aggregate nonlinear model.

This paper examines in depth the issue of individual risk in these systems. The main contributions of the paper are of two kinds:

Risk is modeled and quantified

- (i) The average performance is not an adequate measure of success. It is found empirically that a histogram of QoS is approximately Gaussian, and consequently each load will eventually receive poor service.
- (ii) The variance can be estimated from a refinement of the LTI model that includes a white-noise disturbance; variance is a function of the randomized policy, as well as the power spectral density of the reference signal.

Additional local control can eliminate risk

- (iii) The histogram of QoS is *truncated* through this local control, so that strict bounds on service quality are guaranteed.
- (iv) This has insignificant impact on the grid-level performance, beyond a modest reduction in capacity of ancillary service.

I. INTRODUCTION

The power grid requires regulation to ensure that supply matches demand. Regulation is required by each *balancing authority* (BA) on multiple time-scales, corresponding to the time-scales of volatility of both supply and demand for power. Resources that supply these regulation services are collectively known as *ancillary services*. FERC orders 755 and 745 are intended to provide incentives for the provision of these services. The need for regulation resources increases significantly with greater penetration of renewable generation from wind or sun; this is observed today in regions of the world with significant generation from these sources.

A number of papers have explored the potential for extracting ancillary service through the inherent flexibility of

loads. Examples of loads with sufficient flexibility to provide service to the grid are aluminum manufacturing, data centers, plug-in electric vehicles, heating and ventilation (HVAC), and water pumping for irrigation [1]–[3].

An aggregate of loads can be modeled as a dynamical system. The dynamics of the aggregate depends upon the physical properties of the loads, as well as the mechanism that is used to engage them. For example, with the use of price signals, load response may exhibit threshold behavior in which most loads shut down when the price exceeds some value. Even with direct load control, there may be delay and dynamics, so harnessing ancillary services from flexible loads amounts to a control problem: The BA wishes to design some signal to be broadcast to loads, based on measurements of aggregate power output, so that deviation in power consumption tracks a reference signal – See [4] for detailed discussion.

For many types of loads it has been argued that a randomized control architecture at each load can simplify this control problem [1], [2].




Fig. 1. The control architecture: command ζ is computed at a BA, and transmitted to each load. The control decision at a load is based only on its own state and the signal ζ .

Fig. 1 shows a schematic of the control architecture adopted in [2], [4], in which each load operates according to a randomized policy based on its internal state, and a common control signal ζ . Theoretical results and examples in this prior work demonstrate that local randomized policies can be designed so that control of the loads is easy at the grid level. The analysis was based on a mean field limit (as the number of loads approaches infinity), combined with a linear time-invariant (LTI) system approximation of the aggregate nonlinear model. In the examples considered in this prior work, the linear approximation was found to be minimum phase, which is why simple linear error feedback could be applied to achieve nearly perfect tracking.

Absent in prior work is any detailed analysis of risk for an individual load. In the setting of [2], [4] it can be argued that the quality of service (QoS) for each load is acceptable *on average*, but the volatility of QoS has not been addressed

This research is supported by the NSF grants CPS-0931416 and CPS-1259040, the French National Research Agency grant ANR-12-MONU-0019, and US-Israel BSF Grant 2011506.

Y.C. and S.M. are with the Department of Electrical and Computer Engg. at the University of Florida, Gainesville. A.B. is with Inria and the Computer Science Dept. of École Normale Supérieure, Paris, France.




Fig. 2. Histogram of QoS with and without local opt-out control.

to-date.

This paper examines in depth the issue of individual risk. The main contributions of the paper are of two kinds.

In the first part of the paper we propose approaches to quantify QoS, and methods to approximate its mean and variance so that we can quantify risk. The variance is estimated based on a stochastic LTI model that approximates the dynamics of the time-dependent QoS function. The additive disturbance in the LTI model is a function of the randomized policy used at each load and also the exogenous reference signal. Estimation of the variance of QoS is thus reduced to an estimation of the power spectral density of these disturbances. These theoretical results are developed in Section III.

A simple approach is proposed to restrict QoS to pre-specified bounds: A load will opt-out of service to the grid temporarily, whenever its QoS is outside of pre-specified bounds. This essentially eliminates risk since the histogram of QoS is restricted to these bounds.

Numerical results surveyed in Section IV confirm that the histogram of QoS is truncated through this local control, so that strict bounds on service quality are guaranteed. This has *insignificant* impact on the grid-level performance, beyond a modest reduction in capacity of ancillary service.

Fig. 2 shows a histogram of QoS based on simulation experiments described in Section IV. The plot on the left hand side shows that the Gaussian approximation is a good fit with empirical results when there is no local opt-out control. The figure on the right shows how the histogram is truncated when opt-out control is in place.

We begin with a brief survey of a portion of results from [2], [3], and a precise definition of QoS for a load.

II. RANDOMIZED CONTROL AND MEAN-FIELD MODELS

A. Randomized control

The system architecture considered in this paper is illustrated in Fig. 1, based on the following components:

- (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.
- (ii) Each load evolves as a controlled Markov chain on the finite state space $\mathbf{X} = \{x^1, \dots, x^d\}$. Its transition

probability 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 is given by,

$$\mathbf{P}\{X_{\tau+1}^i = x^j \mid X_\tau^i = x^i, \zeta_\tau = \zeta\} = P_\zeta(x^i, x^j) \quad (1)$$

for each $x^i, x^j \in \mathbf{X}$.

- (iii) 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} , which is also normalized.

An approach to construction of $\{P_\zeta : \zeta \in \mathbb{R}\}$ was proposed in [2], [4] based on information-theoretic arguments. In the present paper we do not require a specific construction, but assume that P_ζ is continuously differentiable in the parameter ζ .

The introduction of the time index t in (i) and τ in (ii) is motivated by the possibility that the sampling time at the grid level may be faster than the sampling time at each load. This is explained in Section II-C.

It is assumed that the power consumption at time t from load i is equal to some function of the state, denoted $\mathcal{U}(X_t^i)$. The normalized power consumption is denoted,

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

The superscript is dropped unless dependency on N must be emphasized.

The nominal behavior of each load is defined as the dynamics with $\zeta \equiv 0$. In this case the loads are independent Markov chains. It is assumed that the Markov chain is ergodic: It is uni-chain and aperiodic, so that P_0 has unique invariant distribution π_0 . The value $y^0 := \sum_x \pi_0(x) \mathcal{U}(x)$ is the average nominal power usage. On combining the ergodic theorem for Markov chains with the Law of Large Numbers for i.i.d. sequences we can conclude that $y_t^N \approx y^0$ when both N and t are large, provided $\zeta \equiv 0$.

It is assumed that the regulation signal is also normalized so that tracking amounts to choosing the signal ζ so that $r_t \approx \tilde{y}_t$ for all t , where $\tilde{y}_t = y_t - y^0$ is the deviation from

nominal behavior. For example, we might use error feedback of the form,

$$\zeta_t = G_c e_t, \quad e_t = r_t - \tilde{y}_t, \quad (2)$$

where G_c is a transfer function. The design of G_c depends on a linear systems approximation for dependency of \mathbf{y} on ζ . This can be obtained by the construction of a mean-field model, as in the prior work [2], [5], [6].

B. Mean-field model

The mean-field model is based on consideration of the empirical distributions,

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

Under very general conditions on the input sequence ζ , it can be shown that the empirical distributions converge to a solution to the nonlinear state space model equations,

$$\mu_{t+1} = \mu_t P_{\zeta_t} \quad (3)$$

The interpretation remains the same: μ_t represents the fraction of loads in each state, and sums to 1. The output is denoted $y_t = \sum_x \mu_t(x) \mathcal{U}(x)$, which is the limit of the $\{y_t^N\}$ as $N \rightarrow \infty$. See [2], [4] for details.

The unique equilibrium with $\zeta \equiv 0$ is $\mu_t \equiv \pi_0$ and $y_t \equiv y^0$. Its linearization around this equilibrium is given by the linear state space model,

$$\begin{aligned} \Phi_{t+1} &= A\Phi_t + B\zeta_t \\ \gamma_t &= C\Phi_t \end{aligned} \quad (4)$$

where $A = P_0^T$, C is a row vector of dimension $d = |\mathcal{X}|$ 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 the matrix \mathcal{E} is equal to the derivative,

$$\mathcal{E} = \left. \frac{d}{d\zeta} P_\zeta \right|_{\zeta=0} \quad (5)$$

In the state equations (4), the state Φ_t is d -dimensional, and $\Phi_t(i)$ is intended to approximate $\mu_t(x^i) - \pi_0(x^i)$, $1 \leq i \leq d$. The output γ_t is an approximation of \tilde{y}_t .

C. Super-sampling

An extension of this model is required in some applications to reduce delay at the grid level. An approach called *super sampling* is introduced without discussion in [2]. Here we give a full description of this approach since it is required in the results that follow.

We let T denote the sampling time for each load, in units of minutes. We fix an integer $m > 1$, and assume that $T_g := T/m$ is the sampling interval at the grid level. The signal ζ is broadcast from the BA every T/m minutes. In applications, T might be equal to 30 minutes, while T/m one minute.

The N loads are assumed to be equally divided into m classes. A load is said to be of class k , where $0 \leq k \leq m-1$, if the reference signal is sampled at times $\{nT + kT/m : n \geq 0\}$. At grid-level time, the state of a load in class k is constant on intervals of length m of the form $\{t = nT + iT/m : k \leq$




Fig. 3. The discrete time control signal $\{\zeta_t\}$ is broadcast at times $t_k = kT/m$, $k = 0, 1, 2, \dots$ in the super-sampled model. A load in class i reads the control signal and makes a decision at consecutive times $\{t_k : k = iT/m + kT, k \geq 0\}$.




Fig. 4. Tracking performance with a collection of “intelligent pools”.

$i \leq k + m - 1$, with potential jumps only at the sampling times $\{nT + kT/m : n \geq 1\}$.

This description is in real time. We henceforth consider a normalized discrete-time model in which $T_g = 1$. The time index $t \in \mathbb{Z}_+$ refers to time at the grid level, and $\tau \in \mathbb{Z}_+$ time for an individual load.

Fig. 3 illustrates the resulting system architecture. A linear approximation is again possible, in which the delay is T/m rather than T .

D. Intelligent pools

The paper [2] on which the present work is based considered an application of this system architecture in which the loads were a homogenous collection of pools. The motivation for considering pools is the inherent flexibility of pool cleaning, and because the total load in a region can be very large. The maximum load is estimated at 1GW in the state of Florida.

In [2] the state space was taken to be the finite set,

$$\mathcal{X} = \{(m, i) : m \in \{\oplus, \ominus\}, i \in \{1, 2, 3, \dots\}\} \quad (6)$$

If $X_t = (\ominus, i)$, this indicates that the load was turned off and has remained off for i time units, and $X_t = (\oplus, i)$ represents the alternative that the load has been operating continuously for exactly i time units. The utility function is equal to the indicator function that a pool is operating, $\mathcal{U}(x) = \sum_i \mathbb{I}\{x = (\oplus, i)\}$.

Fig. 4 shows results from a simulation experiment similar to those surveyed in the prior work [2], [4]. This is based on a symmetric model in which a 12 hour/day cleaning period is desired for each load, and each load consumes 1kW when operating. With 10^5 pools engaged in providing ancillary service to the grid, tracking at the grid level is nearly perfect.




Fig. 5. Histogram of QoS based on simulation of 10^5 pools.

A linear control law of the form (2) was designed based on the mean field model. The behavior of the controlled system is as predicted by the mean-field model.

What was left out in this prior work is any detailed consideration of the quality of service to individual loads. The pool pump example is ideal for illustrating the possibility of risk for an individual load, and illustrating how this risk can be reduced or even eliminated.

Consider the following measure of quality of service (QoS) for an individual load. For some time-horizon T_f , and some function $\ell: \mathcal{X} \rightarrow \mathbb{R}$,

$$\mathcal{L}_\tau = \sum_{i=0}^{T_f} \ell(X_{\tau-i}) \quad (7)$$

In application to the pool pump model we consider the total hours of operation, so that

$$\ell(X_\tau) = \sum_j \mathbb{I}\{X_\tau = (\oplus, j)\} \quad (8)$$

In the same experiment conducted to produce Fig. 4, a histogram of the QoS function (7) was constructed based on the 10^5 pools. The time horizon $T_f = 160$ hours was chosen to approximate one week. The histogram shown in Fig. 5 is based on samples over this period for each of the pools.

The histogram appears Gaussian, with a mean of approximately 80; this is consistent with the time horizon used in this experiment, given the nominal 12 hour cleaning cycle. It is evident that a substantial fraction of pools are over-cleaned or under cleaned by 20 hours or more.

An analysis of QoS is presented in the next section based on a model of an individual load in the mean field limit.

III. MEAN FIELD MODEL FOR AN INDIVIDUAL LOAD

In the mean-field limit, the aggregate dynamics are deterministic, following the discrete-time nonlinear control model (3). The behavior of each load remains probabilistic.

We define the mean field model for *one load* by replacing ζ_t^N with its mean-field limit ζ_t . The justification is that we have a very large number of loads, but our interest is in the statistics of an individual.

Under the conditions under which the mean field model is obtained as a limit, the signal ζ is a function of the regulation signal r that can be expressed as causal feedback,

$$\zeta_t = \phi_t(\mu_0, \dots, \mu_t, r_0, \dots, r_t), \quad t \geq 0,$$




Fig. 6. The input ζ modeled as a stationary stochastic process

where in the mean field model, the sequence of probability measures evolve according to (3). Hence ζ_t is a nonlinear function of r_0, \dots, r_t , and the initial condition μ_0 (which is assumed to be deterministic).

In some of our analysis it is assumed that ζ is a stationary stochastic process. The construction of a stationary model for this input process is based on error feedback of the form (2), where $y_t = \sum_x \mu_t(x) \mathcal{U}(x)$. To approximate the statistics of ζ we take a linear approximation for the dynamics from ζ to y , which is assumed to have the transfer function G_p . It is known that the resulting transfer function from r to ζ is then $G_c / (1 + G_c G_p)$ (see any undergraduate textbook on classical control). Recall that r is obtained by filtering the total regulation signal r^0 . In our approximations, the latter is assumed to be derived from filtered white noise with transfer function G_{wr} .

The super-script i will be dropped in our analysis of a single load. Hence $X(t)$ denotes the non-homogeneous Markov chain whose transition probabilities are defined consistently with (1),

$$\mathbb{P}\{X_{\tau+1} = x^j \mid X_\tau = x^i, r_0, \dots, r_\tau\} = P_{\zeta_\tau}(x^i, x^j) \quad (9)$$

The construction of the mean field model (3) presented in [2] is based on lifting the state space from the d -element set $\mathcal{X} = \{x^1, \dots, x^d\}$, to the d -dimensional simplex \mathcal{S} . For the i^{th} load at time τ , the element $\Gamma_\tau \in \mathcal{S}$ is the degenerate distribution whose mass is concentrated at x if $X_\tau = x$; that is, $\Gamma_\tau = \delta_x$.

We denote by $\{X_\tau^\bullet, \Gamma_\tau^\bullet : -\infty < t < \infty\}$ the two stationary processes obtained with $\zeta_t \equiv 0$ and $X_t \sim \pi_0$ for each t .

A. QoS dynamics and opt-out control

In this paper the local control used to improve QoS is a simple “opt-out” mechanism. We formulate a QoS metric \mathcal{L}_τ similar to (7), along with upper and lower bounds b_- and b_+ . A load will revert to the dynamics P_0 (ignoring the reference signal), whenever $\mathcal{L}_\tau \notin [b_-, b_+]$.

The formulation of QoS as the sum (7) is not easily analyzed. It is more convenient to consider the following discounted sum,

$$\mathcal{L}_\tau = \sum_{k=0}^{\tau} \beta^k \ell(X_{\tau-k}) \quad (10)$$

where the discount factor satisfies $\beta \in (0, 1)$, with $\beta \approx 1$. In simulations presented in Section IV, the discount factor is chosen so that $\beta^n \approx 1/2$ when n corresponds to a one-week time-horizon.

The main goal of this section is to estimate the statistics of $\{\mathcal{L}_\tau\}$ in steady-state. For $\beta \sim 1$, a Gaussian approximation




Fig. 7. Linear model of QoS.

is appropriate, so that it is sufficient to estimate its power spectral density $S_{\mathcal{L}}$.

The power spectral density is estimated through several steps:

- (i) We consider the linearized dynamics, so that (14) is approximated by the linear state space model,

$$\Gamma_{\tau+1} = \Gamma_{\tau} P_0 + D_{\tau+1} \quad (11)$$

- (ii) The disturbance D is the sum of two uncorrelated components, whose power spectral densities are approximated in Proposition 3.1 and Proposition 3.3.
- (iii) The approximation in (ii) is based on a spectral analysis of the input ζ , combined with an analysis of the Markov model obtained with $\zeta \equiv 0$.

- (iv) Lastly, given a trace of data r we obtain an approximation of its power spectral density, and from this we obtain an estimate of the power spectral density of ζ .

Accepting these approximations, the power spectral density for $\{\mathcal{L}_{\tau}\}$ can be obtained via the schematic shown in Fig. 7, in which the row vector $C_{\mathcal{L}}$ is defined so that its i th component is equal to $\ell(x^i)$ for each $1 \leq i \leq d$.

Based on the linear model Fig. 7 and an approximation for the power spectral density of D , we arrive at an estimate of the power spectral density $S_{\mathcal{L}}$. This admits a stable factorization,

$$S_{\mathcal{L}}(\theta) = G_{\mathcal{L}}(e^{j\theta}) G_{\mathcal{L}}(e^{-j\theta})^r,$$

where $G_{\mathcal{L}}$ is a stable spectral factor, which is taken to be a row-vector of transfer functions. On letting $\{g_k\}$ denote its impulse response, $G(z) = \sum_{k=0}^{\infty} g_k z^k$, the variance of QoS is obtained in two forms,

$$\sigma_{\mathcal{L}}^2 = \frac{1}{2\pi} \int_0^{2\pi} S_{\mathcal{L}}(\theta) d\theta = \sum_{k=0}^{\infty} \|g_k\|^2$$

We now proceed with the construction of this linear model.

B. Load dynamics

With this state description, the load evolves according to a random linear system, similar to the deterministic dynamics (3):

$$\Gamma_{\tau+1} = \Gamma_{\tau} G_{\tau+1} \quad (12)$$

in which Γ_{τ} is interpreted as a d -dimensional row vector. The $d \times d$ matrix G_{τ} has entries 0 or 1 only, with $\sum_l G_{\tau}(x^j, x^l) = 1$ for all j . It is conditionally independent of $\{\Gamma_0, \dots, \Gamma_{\tau}\}$, given ζ_{τ} , with

$$\mathbb{E}[G_{\tau+1} | \Gamma_0, \dots, \Gamma_{\tau}, \zeta_{\tau}] = P_{\zeta_{\tau}}. \quad (13)$$

The random linear system (12) can thus be described as a linear system driven by “white noise”:

$$\Gamma_{\tau+1} = \Gamma_{\tau} P_{\zeta_{\tau}} + \Delta_{\tau+1} \quad (14)$$

where, $\{\Delta_{\tau+1} = \Gamma_{\tau}(G_{\tau+1} - P_{\zeta_{\tau}}) : \tau \geq 0\}$ is a martingale difference sequence. We can compute the covariance of the noise in one special case:

Proposition 3.1: Consider the stationary sequence $\{\Delta_{\tau}^{\bullet}\}$ obtained for the system in which $\Gamma_{\tau} = \Gamma_{\tau}^{\bullet}$ for all t . Its covariance is given by,

$$\Sigma^{\Delta} := \text{Cov}(\Gamma_{\tau}^{\bullet}[G_{\tau+1} - P_0]) = \Pi - P_0^T \Pi P_0 \quad (15)$$

where $\Pi = \text{diag}(\pi_0)$. \square

The proof of Proposition 3.1 is contained in the Appendix.

C. LTI approximations

The LTI approximation of the nonlinear system (14) is more complex than the linearized mean field model (4) because of the presence of randomness. We have the small-signal approximation as before: applying (5) gives,

$$\Gamma_{\tau+1} \approx \Gamma_{\tau} P_0 + D_{\tau} \quad (16)$$

where the disturbance consists of two terms,

$$D_{\tau} := B_{\tau}^T \zeta_{\tau} + \Delta_{\tau+1} \quad (17)$$

in which $B_{\tau}^T = \Gamma_{\tau} \mathcal{E}$. The recursion (16) is of the form of the state equation that defines Φ in (4), with the addition of two ‘disturbances’. We can extend Proposition 3.1 to approximate the power spectral density of the overall disturbance, but not in closed form:

Lemma 3.2: Consider the approximate disturbance given by $\Gamma_{\tau}^{\bullet} \mathcal{E} + \Delta_{\tau+1}^{\bullet}$. Its power spectral density can be expressed as the sum,

$$S(\theta) = S_{B\zeta}(\theta) + S_{\Delta}(\theta),$$

where $S_{\Delta}(\theta) = \Sigma^{\Delta}$ for all θ . The power spectral density $S_{B\zeta}(\theta)$ for $\Gamma_{\tau}^{\bullet} \mathcal{E} \zeta_{\tau}$ is the Fourier transform of the product of the autocorrelation functions for $\Gamma_{\tau}^{\bullet} \mathcal{E}$ and ζ_{τ} . \square

Proof: The random vectors $\zeta_{\tau} \Gamma_{\tau}^{\bullet} \mathcal{E}$ and $\Delta_{\tau+1}$ are uncorrelated by construction, which is why we obtain a sum. The sequence $\{\Delta_{\tau+1}\}$ is uncorrelated, so that $S_{\Delta}(\theta)$ is independent of θ . Finally, B_{τ} and ζ_{τ} are independent, which is why we obtain a product of autocorrelation functions. \square

Computation of $S_{B\zeta}$ is illustrated in an example. First we require some additional notation. The functions $\{\varepsilon_k\}$ are defined so that the k th component of B_{τ} is given by $\varepsilon_k(X_{\tau})$:

$$\varepsilon_k(x^j) = \mathcal{E}_{j,k}, \quad 1 \leq j, k \leq d.$$

For any $\varrho \in \mathbb{C}$ satisfying $|\varrho| < 1$, the resolvent matrix is the discounted sum,

$$U_{\varrho} = \sum_{n=0}^{\infty} \varrho^n P^n$$

For two functions $f, g: \mathcal{X} \rightarrow \mathbb{R}$ denote,

$$\langle f, g \rangle = \sum \pi_0(x) f(x) g(x)$$

Proposition 3.3: Consider the special case in which the autocorrelation function for ζ is given by

$$R_{\zeta}(n) = \sigma_{\zeta}^2 \rho^n, \quad n \geq 0, \quad (18)$$




Fig. 8. Scaled BPA signal.

where $\rho \in (0, 1)$ and $\sigma_\zeta^2 > 0$. The i, j entry of the power spectral density matrix is then given by,

$$S_{B\zeta}^{ij}(\theta) = (\langle \varepsilon_i, U_\varrho \varepsilon_j \rangle + \langle \varepsilon_j, U_\varrho \varepsilon_i \rangle^* - \langle \varepsilon_i, \varepsilon_j \rangle) \sigma_\zeta^2 \quad (19)$$

where $\varrho = \rho^m e^{-j\theta}$, and the superscript “*” denotes complex conjugate. \square

The proof of Proposition 3.3 is given in the Appendix.

It is not difficult to extend the result to the case where ζ has a power spectral density, with stable spectral factor G_ζ :

$$S_\zeta(\theta) = G_\zeta(e^{j\theta})G_\zeta(e^{-j\theta}).$$

If the transfer function $G_\zeta(z)$ has distinct poles $\{\rho_1, \dots, \rho_{n_G}\}$, then we can construct constants $\{A_1, \dots, A_{n_G}\}$ such that for $n \geq 0$,

$$R_\zeta(n) = \sum_{j=1}^{n_G} A_j \rho_j^n$$

Based on this representation, the power spectral density $S_{B\zeta}$ can be expressed as a sum of n_G terms, each of the form given in (19).

D. Model for the regulation signal.

This section is concluded with discussion on the construction of the transfer function G_{wr} shown in Fig. 6.

Fig. 8 shows the regulation signal for a typical week within the Bonneville Power Authority (BPA) [7]. The power spectral density plot shown was estimated using Matlab’s `psd` command.

To obtain a statistical model, it is assumed that r^0 evolves as the ARMA model,

$$r_t^0 + a_1 r_{t-1}^0 + a_2 r_{t-2}^0 = w_t^0 + b_1 w_{t-1}^0 \quad (20)$$

in which w^0 is white noise with variance σ_w^2 . The extended least squares (ELS) algorithm was used to estimate the coefficients a_1, a_2, b_1 , and the variance σ_w^2 . Based on data samples in Fig. 8, the ELS algorithm terminated at $[a_1, a_2, b_1]^T = [-0.9009, 0.0365, 0.0859]^T$, and $\sigma_w^2 = 0.005$. In the z -domain, its transfer function is expressed as below,

$$G_{wr}(z) = \frac{1 + 0.08594z^{-1}}{1 - 0.9009z^{-1} + 0.03653z^{-2}} \quad (21)$$

The dashed line in Fig. 8 is the estimate of the spectrum given by $|R^0(e^{jw})|^2 = \sigma_w^2 |G_{wr}(e^{jw})|^2$.

IV. APPLICATION TO INTELLIGENT POOLS

Numerical experiments were conducted on the pool pump model described in Section II-D to investigate the performance of local opt-out control. We show through simulations the improvement on QoS and the impact on the grid with this additional local control. In addition, we provide an analytic formula to estimate QoS.

A. Simulation setup

In order to monitor the health of pools over a long time horizon, a simulated regulation signal was generated using the filter (21). Driven by white noise $N(0, \sigma_w^2)$ with variance $\sigma_w^2 = 0.005$ and sampling time $T_g = 5$ minutes, we obtained a 400-hour long simulated regulation signal r^0 . This signal was then passed through a low-pass filter to generate the reference signal r used in the experiments.

The supersampling approach was used with $m = 6$, corresponding to a local sampling time equal to $T = mT_g = 30$ minutes. The simulation used 10^5 homogeneous Markov models for each pool; each pool pump is assumed to consume 1 KW during operation.

In the QoS function a normalized version of the function (8) is used, defined by,

$$\ell(X_\tau) = \sum_j [\mathbb{I}\{X_\tau = (\oplus, j)\} - \mathbb{I}\{X_\tau = (\ominus, j)\}] \quad (22)$$

If $\mathcal{L}_t^i > 0$ then the pool has received too much cleaning, and if $\mathcal{L}_t^i < 0$ then the pool has been under-cleaned. The acceptable interval for opt-out control was set to $[-20, +20]$ ($b_- = -20$ and $b_+ = 20$, following the notation at the top of Section III-A).

The discount factor in the QoS function (10) was chosen as $\beta = 0.9975$.

B. Risk

The QoS without opt-out control is illustrated in the histogram shown on the left hand side of Figure 2, based on samples of the QoS function (10) from each load over 400 hours. A Gaussian distribution with mean zero and variance σ_ζ^2 is plotted as the dashed line in this figure, based on the variance approximation obtained in Section III. The approximation is close to the simulation result.

C. Performance of opt-out control

The regulation signal used to generate on the left hand side of Fig. 2 was re-used in experiments to investigate the impact of opt-out control on the grid, and on the QoS of an individual load.

QoS is dramatically improved with opt-out control: The histogram shown on the right hand side of Fig. 2 is truncated to the interval $[-20, +20]$ as desired. The dashed line is the conditional Gaussian density, truncated to this interval, with the same mean and variance as the Gaussian distribution shown on the left hand side of this figure. The conditional density is a good fit to the histogram obtained through simulations.

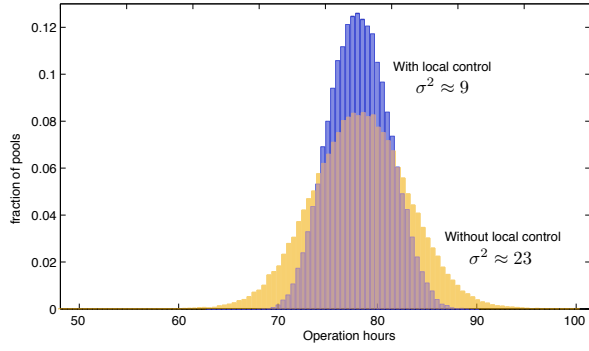


Fig. 9. Comparison of the histograms for the moving window QoS indicator (7), with and without local control.




Fig. 10. Tracking performance with local control.

Fig. 9 shows a comparison of two histograms of the QoS indicator (7), with and without local control. The histogram with larger variance is exactly the same as shown in Fig. 5. The variance is reduced by more than a factor of two using local control.

The grid level tracking performance shown in Fig. 10 remains nearly perfect. This is surprising, given the improvement in QoS. The explanation is that very few loads opt out: In Fig. 11 we see that in this simulation, no more than 3% of loads opt out at any time.

There are limitations on the capacity of ancillary service from a collection of loads, and it is likely that op-out control will reduce capacity. We next obtain an approximation of QoS to better understand the grid-level impact of opt-out control.

D. QoS analysis

The variance of QoS was approximated in Section III based on a small-signal analysis.

The following proposition gives an approximation of its mean at any time. This requires bridging time at the grid and load levels: Any time t can be expressed as $t = m\tau + j$, where m is the super-sampling parameter, $\tau \in \mathbb{Z}_+$, and $j \in \{0, 1, \dots, m-1\}$. The average QoS at time t is denoted,

$$\bar{\mathcal{L}}_t = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_t^i,$$

where \mathcal{L}_t^i is the QoS for load i at time t , defined using grid-level time. For each i , the function $\{\mathcal{L}_t^i : t = 0, 1, \dots\}$




Fig. 11. Percentage of pools that “opt out” as a function of time.




Fig. 12. Tracking performance using a larger reference signal with local control.

is piecewise constant, with potential jumps spaced T/m minutes apart.

A discounted reference sequence is defined by,

$$R_t^\beta = \sum_{k=0}^{\tau} \beta^k r_{m(\tau-k)+j} \quad (23)$$

where N is the number of pools.

Proposition 4.1: Suppose there is perfect tracking: $\tilde{y}_t = r_t$ for all t , where $\tilde{y}_t = y_t - y^0$, and $y^0 = \frac{1}{2}$. Then,

$$\bar{\mathcal{L}}_t = 2R_t^\beta. \quad (24)$$

□

The proof of Proposition 4.1 is given in the Appendix.

In practice \tilde{y}_t only approximates r_t , so (24) is only an approximation. This is illustrated in Fig. 11, where the dotted line is $2R_t^\beta$, and the solid line is $\bar{\mathcal{L}}_t$. They nearly coincide because of the nearly perfect tracking observed in Fig. 10.

If the reference signal r takes a long and large positive or negative excursion, the proposition suggests that many loads will receive poor quality of service, and opt out of service to the grid.

An experiment was conducted to illustrate this conclusion. The reference signal was scaled by 2.2, which is approximately the largest scaling possible for the system without opt-out control. Results are summarized in Fig. 14. In this situation, $2R_t^\beta$ rose over 20 at $t \approx 175hr$, which is the upper bound for an individual load. This caused many loads to opt

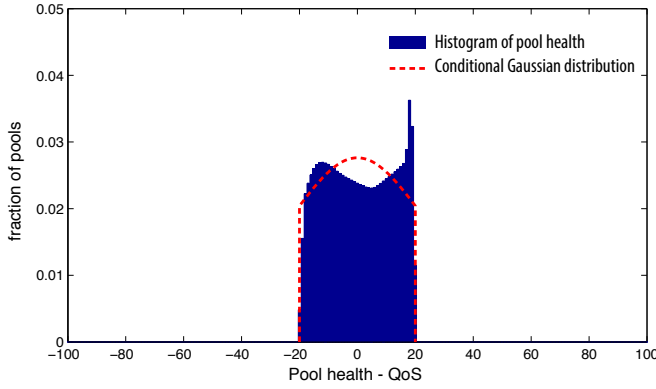


Fig. 13. Histogram of QoS subject to a larger reference signal.

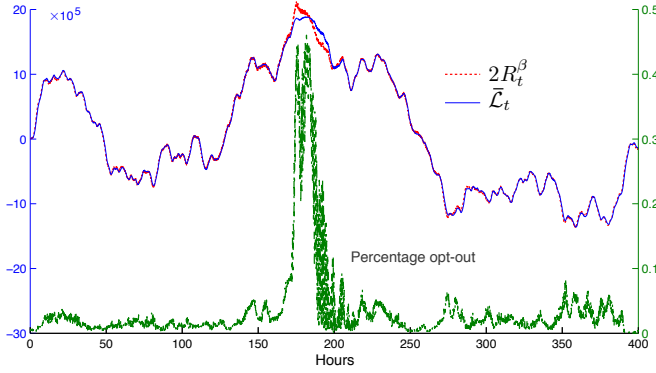


Fig. 14. Percentage of pools that opt out for a larger reference signal.

out, and the corresponding tracking performance degraded at $t \approx 175hr$, as shown in Fig. 12.

The histogram of pool health in Fig. 13 is still within the interval $[-20, +20]$, but it has a greater fraction of pools close to the boundaries ± 20 .

V. CONCLUSIONS

The main technical contribution of this paper is the approximation of QoS for an individual load. It is remarkable that it is possible to obtain accurate estimates of first and second order statistics for an individual load, taking into account second order statistics of exogenous inputs (in this case the reference signal), along with correlation introduced by the Markovian model. It is also remarkable that strict bounds on QoS can be guaranteed while retaining nearly perfect grid-level tracking.

REFERENCES

- [1] 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.
- [2] 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.
- [3] 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.

- [4] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, "Ancillary service to the grid using intelligent deferrable loads," *ArXiv e-prints: arXiv:1402.4600*, 2014, Submitted to the *IEEE Transactions on Auto. Control*.
- [5] 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.
- [6] 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.
- [7] (2013) BPA balancing authority load and total wind, hydro, and thermal generation (website). [Online]. Available: transmission.bpa.gov/business/operations/Wind/baltwg.aspx

APPENDIX

A. Proof of Proposition 3.3.

We set $\sigma_\zeta^2 = 1$ to simplify notation.

The power spectral density matrix $S_{B\zeta}(\theta)$ is the Fourier transform of the autocorrelation matrix for $B\zeta$. The stochastic processes $\mathbf{X} = \mathbf{X}^*$ and ζ are independent under the conditions of the proposition. Consequently, the k, l entry of $S_{B\zeta}(\theta)$ can be represented as,

$$\begin{aligned} S_{B\zeta}^{kl}(\theta) &= \sum_{\tau=-\infty}^{\infty} \mathbb{E}[\varepsilon_k(X_0)\zeta_0][\varepsilon_l(X_\tau)\zeta_{m\tau}]e^{-j\tau\theta} \\ &= \sum_{\tau=-\infty}^{\infty} \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_\tau)]\mathbb{E}[\zeta_0\zeta_{m\tau}]e^{-j\tau\theta} \\ &= \sum_{\tau=-\infty}^{\infty} \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_\tau)]\rho^{|\tau|}e^{-j\tau\theta} \end{aligned}$$

where we have used $\sigma_\zeta^2 = 1$, and each expectation is in steady-state.

The right hand side is expressed as the sum of three terms:

$$\begin{aligned} S_{B\zeta}^{kl}(\theta) &= \sum_{\tau=0}^{\infty} (\rho^m e^{-j\theta})^\tau \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_\tau)] \\ &\quad + \sum_{\tau=-\infty}^0 (\rho^m e^{j\theta})^{-\tau} \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_\tau)] \\ &\quad - \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_0)] \end{aligned} \quad (25)$$

The last term in this expression is expressed as the inner product $\mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_0)] = \langle \varepsilon_k, \varepsilon_l \rangle$.

Using the smoothing property of conditional expectation and the definition $\varrho = \rho^m e^{-j\theta}$, the first term in (25) can be expressed in terms of the resolvent:

$$\begin{aligned} \sum_{\tau=0}^{\infty} (\rho^m e^{-j\theta})^\tau \mathbb{E}[\varepsilon_k(X_0)\varepsilon_l(X_\tau)] &= \sum_{\tau=0}^{\infty} \varrho^\tau \mathbb{E}[\varepsilon_k(X_0)\mathbb{E}[\varepsilon_l(X_\tau)|X_0]] \\ &= \sum_{\tau=0}^{\infty} \varrho^\tau \mathbb{E}[\varepsilon_k(X_0)P^\tau \varepsilon_l(X_0)] \\ &= \mathbb{E}[\varepsilon_k(X_0) \sum_{\tau=0}^{\infty} \varrho^\tau P^\tau \varepsilon_l(X_0)] \\ &= \langle \varepsilon_k, U_\varrho \varepsilon_l \rangle \end{aligned}$$

The second term in (25) is transformed into a form similar to the first term,

$$\begin{aligned} \sum_{\tau=-\infty}^0 (\rho^m e^{j\theta})^{-\tau} \mathbb{E}[\varepsilon_k(X_0) \varepsilon_l(X_\tau)] \\ = \sum_{\tau=0}^{\infty} (\varrho^*)^\tau \mathbb{E}[\varepsilon_l(X_0) \varepsilon_k(X_\tau)] \end{aligned}$$

which is equal to $\langle \varepsilon_l, U_{\varrho} \varepsilon_k \rangle^*$. These representations and (25) establish the desired conclusion (19). \square

B. Proof of Proposition 3.1

The random variables Γ_t and G_{t+1} are independent, with respective means π_0 and P_0 . The row vector Γ_t has entries zero or one, and the same is true for the matrix G_{t+1} .

The k, l entry of the covariance matrix Σ^Δ can be computed as follows. First, using the fact that $\Gamma_t(i)\Gamma_t(j) = 0$ if $i \neq j$,

$$\begin{aligned} \Sigma_{k,l}^\Delta &= \mathbb{E} \left[\left(\sum_i \Gamma_t(i) [G_{t+1}(i, k) - P_0(i, k)] \right) \right. \\ &\quad \left. \left(\sum_j \Gamma_t(j) [G_{t+1}(j, l) - P_0(j, l)] \right) \right] \\ &= \mathbb{E} \left[\sum_i \Gamma_t(i) [G_{t+1}(i, k) - P_0(i, k)] [G_{t+1}(j, l) - P_0(j, l)] \right] \\ &= \sum_i \pi_0(i) \mathbb{E} \left[[G_{t+1}(i, k) - P_0(i, k)] [G_{t+1}(j, l) - P_0(j, l)] \right] \end{aligned}$$

where the last line uses independence of Γ_t and G_{t+1} . Further manipulations give,

$$\begin{aligned} \Sigma_{k,l}^\Delta &= \sum_i \pi_0(i) \left(\mathbb{E} [G_{t+1}(i, k) G_{t+1}(j, l)] \right. \\ &\quad \left. - P_0(i, k) P_0(j, l) \right) \\ &= \sum_i \pi_0(i) \left(\mathbb{E} [G_{t+1}(i, k) G_{t+1}(j, k)] \mathbb{I}\{k = l\} \right. \\ &\quad \left. - P_0(i, k) P_0(j, l) \right) \\ &= \sum_i \left(\pi_0(i) P_0(j, k) \mathbb{I}\{k = l\} - P_0(i, k) P_0(j, l) \right) \end{aligned}$$

where we have used the property that G_{t+1} has Bernoulli entries. The right hand side is precisely (15). \square

C. Proof of Proposition 4.1.

Denote for any time $t = m\tau + j$,

$$\bar{\ell}_t = \frac{1}{N} \sum_{i=1}^N \ell(X_t^i).$$

Based on the definition of y_t we have,

$$\begin{aligned} \bar{\ell}_t &= \frac{1}{N} \sum_{i=1}^N \sum_j [\mathbb{I}\{X_t^i = (\oplus, j)\} - \mathbb{I}\{X_t^i = (\ominus, j)\}] \\ &= \frac{1}{N} \left(2 \sum_{i=1}^N \sum_j \mathbb{I}\{X_t^i = (\oplus, j)\} - N \right) \\ &= 2\tilde{y}_t \end{aligned}$$

Let $\ell_\tau^{C_i}$ denote the average of indicators in class i , and $\mathcal{L}_t^{C_i}$ denote the discounted sum of $\ell_\tau^{C_i}$.

$$\mathcal{L}_t^{C_j} = \sum_{k=0}^{\tau} \beta^k \ell_{\tau+1-k}^{C_j}$$

Note that at time t only $\ell_\tau^{C_j}$ is updated; other $\ell_\tau^{C_i}$, $i \neq j$ remain at their values at $t-1$. Therefore we have $\mathcal{L}_{m\tau+j+n}^{C_j} = \mathcal{L}_{m\tau+j}^{C_j}$, for all j and $n \in \{1, 2, \dots, m-1\}$.

Based on this fact we have,

$$\begin{aligned} \bar{\ell}_t &= \frac{1}{N} \sum_{i=1}^N \ell(X_t^i) \\ &= \frac{1}{N} \left[\frac{N}{m} \left(\sum_{i=j+1}^{m-1} \ell_\tau^{C_i} + \sum_{i=0}^j \ell_{\tau+1}^{C_i} \right) \right] \\ &= \frac{1}{m} \left(\sum_{i=j+1}^{m-1} \ell_\tau^{C_i} + \sum_{i=0}^j \ell_{\tau+1}^{C_i} \right) \end{aligned}$$

Therefore, the average QoS at time t is

$$\begin{aligned} \bar{\mathcal{L}}_t &= \frac{1}{N} \sum_{i=1}^N \mathcal{L}_t^i \\ &= \frac{1}{N} \left[\frac{N}{m} \left(\mathcal{L}_t^{C_1} + \dots + \mathcal{L}_t^{C_m} \right) \right] \\ &= \frac{1}{m} \sum_{k=0}^{\tau} \beta^k \left(\sum_{i=j+1}^{m-1} \ell_{\tau-k}^{C_i} + \sum_{i=0}^j \ell_{\tau+1-k}^{C_i} \right) \\ &= \sum_{k=0}^{\tau} \beta^k \bar{\ell}_{m(\tau-k)+j} \end{aligned}$$

If tracking is perfect as assumed in the proposition, then $\tilde{y}_t = r_t$, or $\bar{\ell}_t = 2r_t$. We obtain from this and (23) the desired conclusion that $\bar{\mathcal{L}}_t = 2R_t^\beta$. \square