Bayesian parameter estimation for direct load control of populations of air conditioners

Nariman Mahdavi [*,**] Cristian Perfumo [*]
Julio H. Braslavsky [*,**]

* CSIRO Energy Technology, PO Box 330, Newcastle 2300, NSW,
Australia. {nariman.mahdavimazdeh, cristian.perfumo,
julio.braslavsky}@csiro.au
** School of Electrical Engineering and Computer Science, The
University of Newcastle, NSW, Australia

Abstract: Recent approaches for direct load control (DLC) of populations of air conditioners (ACs) to provide demand-side services in the electricity grid rely on mathematical models of the aggregate demand dynamics of these populations. These models can be parametrised by the physical characteristics of the ACs in the population, for example their thermal power. The knowledge of how their physical parameters are distributed in the population of real devices is instrumental in the analysis and implementation of controllers based on such models. For large populations, it is typically assumed that these parameters are stochastically distributed according to some probability distribution, e.g., log-normal, which has been effective in simulations. However, the identification of such distribution for a specific population remains an open problem for real-world deployments of DLC. This paper formulates a Bayesian framework for the state and parameter estimation of a previously developed input/output model for the aggregate demand response of heterogeneous populations of ACs. This framework enables us to assign a prior distribution to the parameters of the model, which is then updated using measurements of power demand data for the population to reach a posterior distribution that is more informative about the true value of these parameters. The framework uses sequential Monte Carlo methods, which are well-suited to existing high-performance computer hardware, and aims to provide a way to fill a gap between simulation and real implementation by validating posterior parameter distributions using real measurements. Simulation results indicate that our approach can successfully capture the values defining the distributions of physical parameters in a population simulated by 10,000 ACs with a standard hybrid dynamic model for each device.

1. INTRODUCTION

1.1 Motivation

Air conditioners and other thermostatically-controlled loads (TCLs) possess two main features that make them suitable candidates for demand response. Firstly, altering their behaviour during a short period of time can go unnoticed by the occupants due to the thermal mass in the air conditioned space acting as energy storage. Secondly, these devices can respond rapidly to a control signal, such as "turn off" or "raise temperature" (Callaway, 2009).

When a *population* of ACs is considered as a whole entity to be controlled, additional advantages appear. Reliability is one of them, as a fault in controlling one of many devices is less critical than a fault in controlling, for example, a single large industrial load. Another advantage is the fine granularity achievable in the power output: the compressor of an on-off AC can only be run at its maximum speed or not run at all (i.e., two power levels), whereas when hundreds of these ACs are combined, there are potentially hundreds of different power levels to which the population can be controlled. Additionally, the more ACs there are in

the population, the less control effort is needed to achieve the same desired aggregate power demand, resulting in lower end-user impact.

To be able to control aggregate power demand dynamics, researchers have developed different approaches to mathematically model the population (e.g., Callaway, 2009; Bashash and Fathy, 2011; Perfumo et al., 2012; Mathieu et al., 2013). These models typically depend on the physical characteristics of the ACs, such as thermal power, thermal mass, and level of insulation of the dwelling.

A main challenge in obtaining these mathematical models is the parameter *heterogeneity* present in any real-life population of ACs. There are several sources that contribute to this heterogeneity; for example:

- different AC brands and models will have different thermal power and coefficients of performance,
- different dwellings will have different heat transfer rates and thermal capacitances, and
- different local environment conditions, occupancy patterns and user preferences, which result in different internal heat loads and different desired temperature set-points.

A common and convenient approach to address this challenge is to assume that the physical parameters of the devices are distributed in the population according to a certain probability density function. For example, Callaway (2009) assumes that thermal power, capacitance and resistance are log-normally distributed, while in Mathieu et al. (2013) and Bashash and Fathy (2011) they are assumed to be uniformly or normally distributed.

Perhaps the most accurate method to characterise these distributions for a real-world population is by individually measuring the parameters of each device, namely by performing physical measurements of the dwelling (such as surface area, number and size of windows and building material) as well as logging occupant preferences and usage patterns. While such a thorough study would result in a complete characterisation of every device in the population, it can be expensive and cause privacy concerns amongst participants. Therefore, the nonintrusive identification of the parameters of such distributed models for a real population is a problem of practical interest, and is addressed in the present paper.

1.2 Our approach

We consider the problem of identifying, from experimental data, probabilistic models to characterise how AC physical parameters are distributed in a population. The experimental data of interest is the aggregate demand of the population (mixed or not with demand from other, non-controllable loads), preferably obtained during periods when dynamic transients occur, such as after a simultaneous change in thermostat set-points.

The parameters we seek to estimate are those of probabilistic distributions, such as means and variances of probability density functions. We will refer to these as *hyper-parameters* (of distributions) to differentiate them from the *physical parameters* (thermal power, thermal resistance and thermal capacitance) of individual devices.

To model the aggregate dynamics of the population we use an input/output linear time-invariant (LTI) model introduced in Perfumo et al. (2012). As shown in Perfumo et al. (2012) and Braslavsky et al. (2013), this model can capture the dominant dynamics in aggregate demand response with sufficient accuracy and robustness to allow tight model-based load control.

While the model from Perfumo et al. (2012) has a simple dynamic structure, it is parametrized by complex (but known) nonlinear functions of the distribution hyper-parameters that we seek to estimate, which then requires nonlinear identification techniques. While one could employ approximate methods, such as the Extended Kalman filter (EKF), the intrinsic approximate nature of an approach based on linearisation may present convergence issues even for moderate nonlinearities in these models (Mathieu et al., 2013).

In this paper we adopt a Bayesian approach, which conveniently combines prior knowledge and experimental data to infer information about the system (Schön, 2003; Gordon et al., 1993). The Bayesian approach has received a great deal of attention in the parameter estimation literature (e.g., Jeffrey, 2004; Doucet et al., 2001), however,

it has not been more broadly used mainly due to the computational difficulties involved in computing posterior densities. Fortunately, there has been significant progress in recent years both in the Bayesian inference theory and its implementation on high-performance computing hardware. A breakthrough has been the introduction of Markov Chain Monte Carlo methods. From this family of methods, Sequential Monte Carlo (SMC) methods are specially suited to joint parameter and state estimation of complex nonlinear systems, successfully demonstrated in real-world applications in recent years (Dowd, 2011).

We apply the Bayesian approach by first postulating prior distributions for the population hyper-parameters. From these prior distributions we obtain a parametrization for the LTI model and a prior aggregate dynamic response. Subsequently, this response and the observed aggregate demand are used to update the prior distributions to obtain posterior distributions. The means of these posterior distributions give estimates of the means and variances of the physical parameters in a population.

The approach is validated through simulations, using observed data generated with "perfect" knowledge of the hyper-parameters and empirical distributions of the physical parameters. We simulate a population of 10,000 ACs operating independently and subject to a simultaneous step in thermostat set-points. The Bayesian inference engine is implemented using *LibBi* (Murray, 2013), a versatile computational environment for Bayesian inference.

The simulation results show that the proposed Bayesian approach can produce estimates that are close to the real values even when the prior distributions are wide and uninformative. The method complements the parametric second order LTI model providing a complete toolkit for practical model-based direct load control of TCLs.

We describe in more detail the identification problem considered in Section 2 . Section 3 provides a brief revision of the Bayesian inference framework and SMC methods used. The simulation scenarios used to test our approach are presented in Section 4 describes, followed by a discussion of the results obtained in Section 5.

2. PROBLEM STATEMENT

The Bayesian-based identification problem considered requires two main components: a parametric model of the AC population aggregate demand response, and experimental data of such response from a population.

2.1 Parametric model used for Bayesian estimation

We consider the LTI second-order system proposed in Perfumo et al. (2012) to model the aggregate power demand response $y^p(t)$ of a large heterogeneous population of ACs, each of which is modelled as a relay-driven dynamic system (Chong and Debs, 1979) regulating around a temperature set-point. According to this model, the aggregate response to a simultaneous step change in set-points is given by

$$y^p(t) = D_{ss}(T_r) - D(t), \tag{1}$$

$$\dot{x}(t) = \begin{bmatrix} -2\xi\omega_n & -\omega_n^2 \\ 1 & 0 \end{bmatrix} x(t) + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u(t) \tag{2}$$

$$D(t) = \begin{bmatrix} b_1 - 2\xi\omega_n b_2 & b_0 - b_2\omega_n^2 \end{bmatrix} x(t) + b_2 u(t) \tag{3}$$

where $D_{ss}(T_r)$ is the (mean) asymptotic steady-state aggregate demand of the population with constant temperature set-points T_r, $D(t)$ is the output of a linear system, $x(t)$ is the state vector and $u(t)$ is a control signal in the form of a common temperature set point offset.

The derivation of the linear model (1)-(3) is based on the assumption that all of the ACs in the population have the same thermal resistance R and thermal power P, and only the thermal capacitance C is log-normally distributed with mean μ_C and relative standard deviation $\sigma_{rel} = \sigma_C/\mu_C$.

The model coefficients b_2, b_1, b_0, ω_n and ξ in (2)-(3) are given as explicit functions of the hyper-parameters μ_C and σ_{rel} as follows (see Perfumo et al. (2012); Perfumo (2013) for the derivation of these values):

$$\omega_n = \frac{\pi\mu_v}{(\sqrt{6}-a)\sqrt{1-\xi^2}}, \quad \xi = \frac{\log(r)}{\sqrt{\pi^2 + \log^2(r)}}, \tag{4}$$

$$b_0 = \frac{\omega_n^2(D_{ss}(T_r) - D_{ss}(T_r + 0.5))}{0.5}, \quad b_2 = D_{ss}(T_r),$$

$$\text{and} \quad b_1 = \omega_n\Bigg[\xi(3D_{ss}(T_r) - 2D_{ss}(T_r + 0.5))+$$

$$\cot(t^1\omega_n\sqrt{1-\xi^2})\sqrt{1-\xi^2}(D_{ss}(T_r) - 2D_{ss}(T_r + 0.5))+$$

$$\frac{2e^{t^1\omega_n\xi}\sqrt{1-\xi^2}(D_{ss}(T_r + 0.5) - D^1)}{\sin(t^1\omega_n\sqrt{1-\xi^2})}\Bigg], \tag{5}$$

where

$$t^1 = (a - \tfrac{7}{8})/\mu_v; \quad D^1 = \frac{1}{6} + \frac{1}{6}\text{erf}\left[\frac{-\log(\mu_v t^1 + \tfrac{7}{8})}{\sqrt{2}\sigma_{rel}}\right]$$

$$a = \exp\left[\log(\sqrt{2}) - \sigma_{rel}^2\tfrac{\log(3)}{\log(2)}\right]; \mu_v = \frac{(T_a - T_r)(1 + \sigma_{rel}^2)}{R\mu_C},$$

$$r = \left|\frac{\text{erf}\left[(0.9 + \sqrt{8}\sigma_{rel})^{-1}\right] - \tfrac{1}{2}}{\text{erf}[0.9^{-1}] - \tfrac{1}{2}}\right|; \quad \sigma_{rel} = \sigma_C/\mu_C \tag{6}$$

and

$$D_{ss}(T) = \left(1 + \frac{\log(1 + \frac{H}{T_a - \theta - H/2})}{\log(1 + \frac{H}{PR + T - T_a - H/2})}\right)^{-1}. \tag{7}$$

Here, erf[.] is the Gauss error function and H is the ACs' thermostat hysteresis deadband width.

2.2 Observations used for Bayesian estimation

In practice, the observation data would be power measurements from the real population of ACs whose parameters we are trying to identify. In order to assess the performance of the proposed approach under perfect population knowledge conditions, we generate virtual observations by simulating 10,000 ACs using a standard hybrid dynamic model for each device (Chong and Debs, 1979; Ihara and Schweppe, 1981) with parameters distributed by suitable stochastic distributions.

The hybrid dynamic model for each AC in the population gives the temperature T_i and thermostat relay state m_i

(the subindex i denotes the ith AC in the population, $i = 1, 2, \ldots, 10,000$) as

$$\frac{dT_i(t)}{dt} = -\frac{1}{C_i R_i}[T_i(t) - T_a + m_i(t)R_i P_i], \tag{8}$$

with the relay hysteresis control

$$m_i(t^+) = \begin{cases} 0 & \text{if } T_i(t) \leq 19.5 + u(t) \\ 1 & \text{if } T_i(t) \geq 20.5 + u(t) \\ m_i(t) & \text{otherwise}, \end{cases} \tag{9}$$

where $u(t)$ is a simultaneous 0.5°C step-change to the temperature set-points of all ACs.

The aggregated normalized power demand y_t (i.e., our observations) is then given by

$$y_t = \frac{\sum_{i=1}^n m_i(t)\frac{P_i}{\text{COP}_i}}{\sum_{i=1}^n \frac{P_i}{\text{COP}_i}}, \tag{10}$$

where COP_i is the coefficient of performance of the ith AC, assumed constant as is common in the literature.

To account for heterogeneity in this ideal population, we distribute the thermal parameters R, P and C using a log-normal density function (Callaway, 2009), with distribution parameter values as listed in Table 1.

Table 1. Simulation parameters.

Parameter	Value	Description
μ_R	2 °C/kW	Mean thermal resistance
μ_C	216 kWmin/°C	Mean thermal capacitance
μ_P	6 kW	Mean thermal power
T_r	20 °C	Temperature set-point
T_a	26 °C	Ambient temperature
σ_{rel}	0.2	Ratio of Standard deviation to the mean value of log-normal distributions for R, C and P

With above mentioned dynamical equations of the ideal population, 800 one minute sampled data (about 13 hours) are generated after half a Celsius degree simultaneous change is applied in thermostat set-points of all devices at 100^{th} time step. The benefit of these artificially generated observations is that we know the true mean and variance of the parameters R, P and C in the population. This will allow us to precisely quantify how close the estimates are from the real values, which would be very difficult using actual devices. In the remainder of the paper we focus on obtaining these values only using the model (1) and the observations generated by (10).

3. BAYESIAN FRAMEWORK

In this section, we formulate our problem in the Bayesian framework. This formulation is concerned with the joint state and parameter estimation of stochastic dynamical systems, which is introduced next.

3.1 Stochastic model

The stochastic extension of (2) could be easily obtained by considering an additive white Gaussian noise. The use of stochasticity may be interpreted as representing uncertainty in the formulation of the model. In addition, since the observations are only available as discrete-time

samples, we consider an Euler-discretization of the model (1)-(3), namely

$$y_t^p = D_{ss}(T_r) - D_t, \qquad (11)$$

$$x_t = (I_2 + A\Delta)x_{t-1} + \begin{bmatrix} \Delta \\ 0 \end{bmatrix}(u_t + w_t), \qquad (12)$$

$$D_t = Cx_t + b_2 u_t, \qquad (13)$$

where I_2 is the 2×2 identity matrix, $\Delta = 1$ minute is the sample time and $w_t \sim \text{Normal}(0, 0.3)$ is the white Gaussian noise.

The state equation (12) represents a Markovian transition equation which moves the system state forward and depends only on the previous state. The Markovian property of equation (12) implies that the state at any time can also be written as the distribution

$$x_t \sim p(x_t|x_{t-1}, \theta), \qquad (14)$$

where θ is the hyper-parameter vector (for example, $\theta = [\sigma_{rel}\ \mu_C]$ or $\theta = [\sigma_{rel}\ \mu_C\ \mu_R\ \mu_P]$, depending on the scenarios which will define in the next section). The prior distribution over the initial state is given by

$$x_0 \sim \text{Uniform}(0, 0.1).$$

Equation (12) can be used to generate realizations of the stochastic process. The ensemble properties of these realizations can be described by $p(X_t|x_0, \theta)$, where X_t is the state over time, i.e., $\{x_i\}_{i=0}^t$.

To be able to relate the output y_t^p of the linear model (11) to the observations y_t in (10) we model the measurement noise as

$$y_t \sim \text{Normal}(y_t^p, 0.2). \qquad (15)$$

This observation model could be considered as the conditional distribution

$$y_t \sim p(y_t|x_t, \theta). \qquad (16)$$

The observation set is denoted by Y_t, which is the measurements up to time t, i.e., $\{y_i\}_{i=0}^t$.

Available measurement information is then used to estimate X_t as well as the hyper-parameter θ. Note that, since the states of the system are unknown due to the Gaussian noise w_t in (12), we have to estimate both the states and the hyper-parameters using the SMC method.

3.2 Recursive estimation

A complete solution for the above mentioned stochastic system at any time t is given by $p(x_t, \theta|Y_t)$, which is the joint probability density function of the state, x_t, and parameters, θ, given all available observations, Y_t. This joint distribution could be written as

$$p(x_t, \theta|Y_t) = p(\theta|Y_t)p(x_t|Y_t, \theta), \qquad (17)$$

where the first factor on the right hand side represents *parameter estimation*, and the second *state estimation*.

In state estimation, we consider θ is known to estimate the state x_t. This is the classic filtering problem, and its solution is given by (see, for example, Schön, 2003),

$$p(x_t|Y_t, \theta) \propto p(y_t|x_t, \theta)p(x_t|Y_{t-1}, \theta), \qquad (18)$$

$$p(x_t|Y_{t-1}, \theta) = \int p(x_t|x_{t-1}, \theta)p(x_{t-1}|Y_{t-1}, \theta)dx_{t-1}. \quad (19)$$

Equations (18) and (19) together provide a recursive scheme to update the filter density as new measurements arrive, i.e., a single stage transition from $p(x_{t-1}|Y_{t-1}, \theta)$ to $p(x_t|Y_t, \theta)$. In other words, by recursive estimation, we mean an estimate of the current state x_t, given information about the last estimate, parameters and the current measurement of y_t.

In parameter estimation, due to Bayes' Theorem, we have

$$p(\theta|Y_t) \propto p(Y_t|\theta)p(\theta), \qquad (20)$$

where the posterior distribution of the parameters $p(\theta|Y_t)$ depends on a likelihood $p(Y_t|\theta)$ and a prior $p(\theta)$. In studying an AC population, if some information is available about the physical parameters of a population, we can use this knowledge to start with a more informative prior distribution. Otherwise, the prior distribution may be chosen as uniform over a suitable range of parameter values.

3.3 Sequential Monte Carlo methods

The main idea in SMC methods is to use a set of random samples, with associated weights, to represent the posterior density (17), (Schön, 2003). To do this, we have to estimate the hyper-parameters and the states, which is done with a particle filter to compute (19),(20), and hence produce a sample from $p(x_t|Y_t, \theta)$ and $p(\theta|Y_t)$. The particle filter is a member of the family of SMC methods, first introduced by Gordon et al. (1993) and applied successfully since in many domains (Doucet et al., 2001).

To implement this particle filter, we have used the software package *LibBi*, which is specially formulated for Bayesian state-space modelling on high-performance hardware (Murray, 2013).

4. SCENARIOS FOR BAYESIAN FORMULATION

We consider three different scenarios to test the proposed Bayesian inference problem and analyse the effect of increased stochasticity and different prior knowledge assumptions on the hyper-parameters. These scenarios and the corresponding results (Section 5) are summarised in Table 2.

Since in the model (1)-(3) only the thermal capacitance C is distributed (log-normally) in the population, we first consider the hyper-parameter vector $\theta = [\sigma_{rel}\ \mu_C]$ in Scenario 1. Subsequently, in Scenario 2, we extend the analysis to the case where R and P are also distributed in the population, i.e., $\theta = [\sigma_{rel}\ \mu_C\ \mu_R\ \mu_P]$. This extension increases the stochasticity of the system. Note that the uniform distribution used as prior knowledge of the hyper-parameters, in both Scenarios 1 and 2, is the least informative distribution for the Bayesian estimation and identification problem. The boundaries of the uniform distributions are chosen around the nominal values of the population as listed on Table 1.

The normal distribution of Scenario 3 represents a prior knowledge about the mean thermal power of the ACs in the population. This information could usually be obtained by direct measurements and is given in terms of the mean and the variances of the physical parameters. Therefore, instead of a uniform prior distribution for μ_P in Scenario 2, we assign a normal distribution to it.

Table 2. Scenarios and results

	Identified hyper-parameters	Prior knowledge	Mean values of the hyper-parameters' posterior distribution; $\theta_{actual} = \begin{bmatrix} 0.2 & 216 & 2 & 6 \end{bmatrix}$	MSE between mean of prior/ posterior distribution y_t^p and observation y_t	
				Prior	Posterior
Scenario 1	$\theta = \begin{bmatrix} \sigma_{rel} & \mu_C \end{bmatrix}$	$\sigma_{rel} \sim$ Uniform $(0.05, 0.3)$ $\mu_C \sim$ Uniform $(180, 250)$	$\hat{\theta} = \begin{bmatrix} 0.23 & 238 \end{bmatrix}$	5.7e-3	5.85e-4
Scenario 2	$\theta = \begin{bmatrix} \sigma_{rel} & \mu_C & \mu_R & \mu_P \end{bmatrix}$	σ_{rel}, μ_C Same as Scenario 1 $\mu_R \sim$ Uniform $(1, 3)$ $\mu_P \sim$ Uniform $(5, 7)$	$\hat{\theta} = \begin{bmatrix} 0.2 & 237.7 & 2.2 & 6.3 \end{bmatrix}$	5.02e-3	4.07e-4
Scenario 3	Same as Scenario 2	$\sigma_{rel}, \mu_C, \mu_R$ Same as Scenario 2 $\mu_P \sim$ Normal $(6, 0.5)$	$\hat{\theta} = \begin{bmatrix} 0.2 & 214.5 & 2.45 & 5.8 \end{bmatrix}$	5.13e-3	8.72e-4

Note that *we do not assume any specific distribution for the physical parameters of the population in the Bayesian estimation process.* We only consider their means and variances as a hyper-parameter and find the posterior distribution of the hyper-parameter (not the physical parameter) based on the observation set. One can then calculate the mean of the hyper-parameter posterior distribution to reach to an estimation for the empirical mean/variance of the population distribution of physical parameters.

5. RESULTS

We now present the results from applying the Bayesian inference approach described in Section 3 to identify the hyper-parameters that govern the aggregate power response of a population of ACs. All the simulations have been done in *LibBi* (Murray, 2013), running on CSIRO's supercomputer cluster Bragg.

We start with the uninformative prior distributions described for Scenarios 1, 2 and 3 presented in Section 4, and aim to obtain more informative posterior distributions. Ideally, these posterior distributions should be centered on the real values of the distributed physical parameters. However, due to possible model errors, there may be some discrepancy.

Figure 1 shows the histograms of the posterior samples of hyper-parameters in Scenario 1. These histograms should not be mistaken with the actual distributions of the physical parameters in the population. The intended meaning of these histograms is "the likelihood of the hyper-parameters (i.e., mean, variance) taking a certain value". We observe that even starting with a very uninformative (uniform) prior distributions, the obtained posterior distributions have much less variance and, thus, are more descriptive. Moreover, the means of these obtained posterior distributions for μ_C and σ_{rel} (238 and 0.23 respectively) are close to the true values (216 and 0.2) listed in Table 1. Note that the results for all scenarios are summarised in Table 2.

Figure 2 shows the prior and posterior distribution of the aggregated power demand $y^p(t)$ in Scenario 1. It can be observed that while the 95% confidence interval of the prior distribution is very wide, the posterior distribution is very narrow around the observations. Additionally, according to the Mean Squared Error (MSE) index listed in Table 2, the mean of the posterior distribution is much closer to the observations than the mean of the prior.

With the same approach to the first scenario, the histograms of the posterior samples of hyper-parameters and

prior/posterior distribution of the aggregated power demand $y^p(t)$ in Scenarios 2 and 3 are obtained and their mean values are listed in Table 2. However, unlike the first scenario, Scenarios 2 and 3 assume that the thermal resistance R and thermal power P are also distributed in the population, which represents a more realistic situation. We can see in Table 2 that the posterior distributions of μ_R and μ_P of Scenario 2 converge to values that are very close to the real one, and that the obtained mean for μ_C is slightly better than that obtained in Scenario 1. Furthermore, the estimated mean value of $\sigma_{rel} = 0.2$ is a much closer estimation than that obtained with Scenario 1. We attribute this improvement to the extra flexibility introduced by the increased degree of stochasticity of Scenario 2 over Scenario 1 (the former has two more random variables than the latter).

Let us now analyse the posterior samples of hyper-parameters in Scenario 3, which uses a more descriptive prior distribution for μ_P. The expected result of having extra information about the physical parameters is that the mean of the resulting posteriors distributions are closer to the true value. However, we can see on Table 2 that we obtained mixed results for Scenario 3. On the one hand, the estimate of μ_C is much closer to the real value than with the previous scenarios, and differs from the true value by less than one percent. Additionally, the estimated value of μ_P is slightly better in the case of scenario three. On the other hand, the obtained estimate for μ_R is further away from the real value than the estimate obtained in Scenario 2. Further work is needed to determine the reason for the degradation in the estimate of μ_R, but one possible explanation could be that the normal prior distribution that we chose had too high variance, thus not being substantially more informative than the uniform distribution used in the previous scenarios.

A comparison between the prior and posterior distributions of the aggregate power demand $y^p(t)$ in Scenarios 2 and 3 can be seen in in Table 2 together with the MSE between their mean and observations. We see that the mean of the posterior distribution is about 10 times closer to the observations than the prior distribution in all Scenarios. Note that extra degrees of freedom of Scenario 2 lead to a much wider prior distribution, because there is more uncertainty. However, the obtained posterior distribution follows the observations with less MSE than that from Scenario 1.
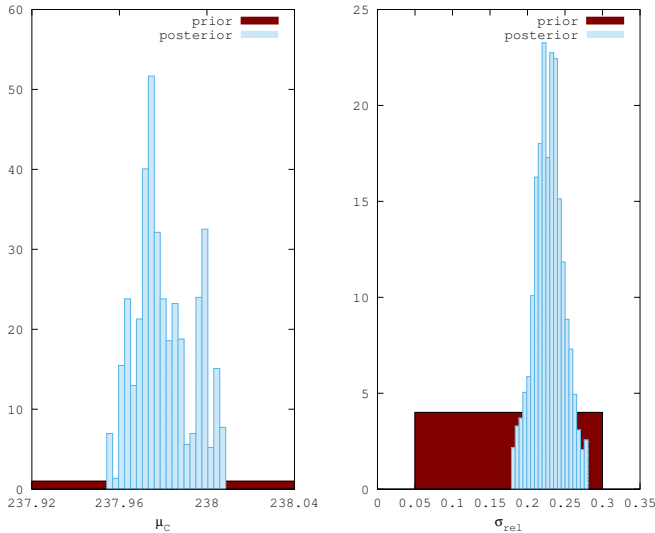
Fig. 1. Histograms of the posterior samples of hyper-parameters μ_C (left) and σ_{rel} (right) in Scenario 1 against their prior distributions.
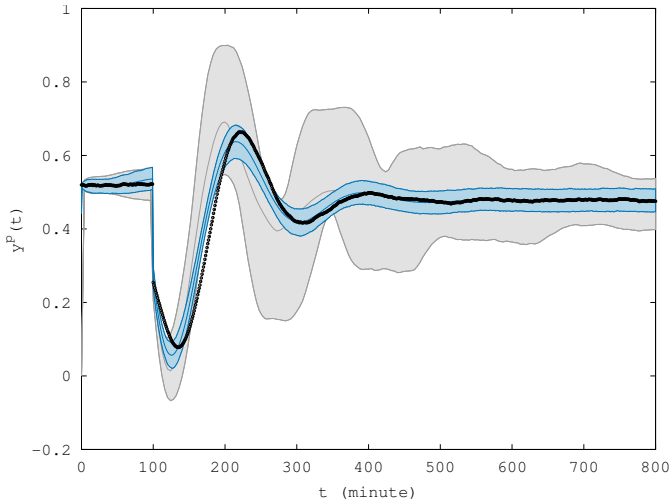

Fig. 2. Scenario 1, prior (gray) vs. posterior (dark, blue) distribution of the aggregated power demand $y^p(t)$ in (1). The bold lines represent the means, and the shaded regions the corresponding 95% confidence intervals at each time. The observations y_t in (10) are plotted in bold black line.

6. CONCLUSIONS

In this paper we have applied Bayesian inference to estimate the physical parameters of a population of ACs. According to the simulation results, our approach results in estimates that are close to the real value even when the prior distributions we start with are wide and uninformative and using a simple second order LTI model to characterise the population of ACs. Thus, a Bayesian-based approach such as the one presented in the paper appears a promising candidate to bridge the current gap between population models that require a probability distribution for the parameters and real groups of ACs. Certainly, it is a much more convenient alternative to individually surveying the dwellings to characterise these distributions.

REFERENCES

S. Bashash and H. K. Fathy. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In *Proc. American Control Conference*, july 2011.

J. H. Braslavsky, C. Perfumo, and J. K. Ward. Model-based feedback control of distributed air-conditioning loads for fast demand-side ancillary services. In *Proc. of the 52nd IEEE Conference on Decision and Control*, Florence, Italy, December 2013. To appear.

D. S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(5):1389–1400, 2009.

C. Chong and A. Debs. Statistical synthesis of power system functional load models. In *18th IEEE Conference on Decision and Control*, volume 18, pages 264 –269, December 1979.

A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo methods in practice.* Springer, 1st edition, 2001.

M. Dowd. Estimating parameters for a stochastic dynamic marine ecological system. *Environmetrics*, 22(4):501–515, 2011.

N. Gordon, D. Salmond, and A. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings on radar and signal processing*, 140:107–113, 1993.

S. Ihara and F.C. Schweppe. Physically based modeling of cold load pickup. *Power Apparatus and Systems, IEEE Transactions on*, PAS-100(9):4142 –4150, September 1981.

R. Jeffrey. *Subjective probability: The real thing.* Cambridge University Press, 1st edition, 2004.

J.L. Mathieu, S. Koch, and D.S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems*, 28(1):430–440, 2013.

L. M. Murray. Bayesian state-space modelling on high-performance hardware using LibBi. *arXiv:1306.3277 [stat.CO]*, 2013.

C. Perfumo. *Dynamic modelling and control of heterogeneous populations of thermostatically controlled loads.* PhD thesis, Electrical Engineering and Computer Science, The University of Newcastle, Callaghan, NSW, Australia, 2013.

C. Perfumo, E. Kofman, J. H. Braslavsky, and J. K. Ward. Load management: Model-based control of aggregate power for populations of thermostatically controlled loads. *Energy Conversion and Management*, 55:36–48, 2012.

T. Schön. *On computational methods for nonlinear estimation.* Linköping Studies in science and technology, Thesis No. 1047, Linköping, Sweden, 2003.