Perfumo, Cristian; Braslavsky, Julio H.; Ward, John K "Model-based estimation of energy savings in load control events for thermostatically controlled loads". Published in IEEE Transactions on Smart Grid Vol. 5, Issue 3, p. 1410-1420 (2014)

Available from: http://dx.doi.org/10.1109/TSG.2014.2298840

Model-based estimation of energy savings in load control events for thermostatically controlled loads

Cristian Perfumo, Julio H. Braslavsky and John K. Ward

Abstract—**Load control (LC) of populations of air conditioners (ACs) is considered suitable to shift energy from on- to off-peak times, and track the intermittent power output of renewable generation. From a technical and economical point of view, it is paramount to quantify the amount of energy that can be saved by implementing these LC events. This paper proposes a new causal methodology to estimate such energy savings using a Kalman filter that includes a parametric second-order model of the aggregate demand of a population of ACs. The proposed methodology relies only on readings of aggregate electrical power at the feeder level and does not require historical load data, or a control group, and hence, it can be used where other methods reported in the literature are inapplicable. The proposed estimator is evaluated on a numerical case study that embeds simulated ACs in real power and temperature data from a 70-house residential precinct.**

I. Introduction

Balanced generation and demand is paramount for the operation of any electrical power network and is traditionally achieved by adjusting the supply to the demand. Lately, however, instead of building more supply infrastructure, there is an increasing interest in solving these problems from the demand side, making a more efficient use of already existing resources [1], [2]. Additionally, the interest in demand-side management will increase as the *smart grid* evolves towards a collection of so-called energy hubs where generation, storage and consumption are coupled [3].

Air conditioners (ACs) and other thermostatically controlled loads (TCLs) such as fridges and space and water heaters have great potential for demand-side services because of their rapid response and thermal inertia (which reduces impact to end-users). In fact, trials throughout the world demonstrate the benefits of externally controlling these types of loads [4], [5], [6], [7], [8], [9].

In recent years, the load control (LC) of large, widely distributed populations of TCLs has gained importance [10], [11], [12], [9], [13]. LC comprises the direct manipulation of the power demand of the electrical devices in these populations to achieve a desired collective response. For example, the temperature set point of the ACs can be controlled by an aggregator to shift energy from on to off-peak times or to provide ancillary services in the energy market.

Cristian Perfumo (cristian.perfumo@csiro.au), Julio H. Braslavsky (julio.braslavsky@csiro.au) and John K. Ward (john.k.ward@csiro.au) are with the Division of Energy Technology at the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia. Julio Braslavsky is also with the School of Electrical and Computer Engineering at the University of Newcastle, Australia.

ACs are especially interesting for LC for two reasons. Firstly, they can provide fast responses with minimal end-use disruptions. In fact, LC of TCLs has been advocated as a demand-side alternative to ancillary services traditionally provided by controlling generation [14], [7], [2], [10], [1]. Secondly, ACs are a prominent driver of peak electricity demand in high temperature days. In this paper we focus on residential ACs, although our approach could easily be applied to other types of ACs as well as other TCLs such as heaters or cool rooms.

It has been shown in the literature that ACs can be controlled successfully to shape their aggregate power demand [15], [16], [13]. However, for LC of ACs to become common in the electricity grid, it is crucial to be able to accurately quantify the energy shifted by a LC event [4], [17]. This quantification helps to assess the financial benefit of LC and define the settlements offered to the AC owners [18].

Unfortunately, the amount of energy shifted by a LC event can only be estimated, as finding the exact figure implies knowing what the consumption would have been if the LC event had not taken place (this constitutes another example of the difficulty of calculating smart-grid related costs [19]). Adopting the terminology in [18] we will refer to the amount of energy shifted as “energy reductions” or “savings” because commonly load shaping using LC is targeted at reducing the amount of energy used during the event (even though there might not be net savings when considering possible “cold-load pick-up” situations after the event).

The existing approaches to quantifying such energy reductions (both in the literature and those used by electricity utilities) can be arranged in three categories: based on historical data, based on system identification, and based on a control group.

Estimates based on historical data can be calculated using power averages recorded during similar-temperature non-LC days or during a certain number of non-LC days prior to the LC event [18], [17]. These methods, which have been used by utilities to quantify savings, are amongst the ones resulting in the worst estimates [18], [17].

Methods based on system identification correlate ambient temperature, time of the day and other data (inputs) to power demand (output) [18], [17], [9]. In particular, regression methods yield good estimates but may require large amounts of historical data to identify the coefficients successfully [18]. Extensive data is also required for other model identification techniques such as polynomial fitting [20] or artificial neural network training [21].

Alternatively, a population of ACs of similar characteristics can be used used as a control group [17], [18]. This approach presents good results as long as the control group is properly selected and both the control and the LC population are large (thousands of devices according to [18]).

The main contribution of this paper is a Kalman-filter based methodology to estimate energy savings during load control events without resorting to historical data or the use of control groups. In particular, the proposed methodology is especially relevant in load control scenarios for ACs connected to the same electrical distribution feeder, for which:

- it is difficult to find a representative control group; or
- it is impractical to estimate uncontrolled demand using existing methods, due for example to the lack of sufficient logged data to produce reliable load forecasts, or due to the higher variability of the load associated with smaller populations.

In this paper we estimate LC energy savings using three different causal approaches that do not require forecast or estimation of non-LC load. The first two approaches are simple rolling fits of measured load averages and ramp rates used in practice (e.g., [22]), and are considered here for comparison purposes. The third approach is the proposed methodology (the main contribution of this paper), which is based on a Kalman Filter designed using a reduced order parametric model for aggregate power demand of a population of ACs. This parametric model was developed in [23] and is briefly revisited in Section II below.

The three approaches to estimate energy savings are evaluated on a series of LC events performed on a small population of ACs, which is simulated using real load and temperature data from a small residential precinct. The LC events considered range from 1 to 4 hours in duration, and are implemented using the integral-action controller designed in [23]. The Kalman filter estimates are shown to consistently outperform the simpler estimates based on constant rolling demand averages or constant demand ramp rates. Note that integrating a model of a population of ACs into a Kalman filter has been proposed in [24], but assuming the availability of a utility load forecast with white noise forecast error. Our approach can be applied when such forecast is not available.

The remainder of this paper is organised as follows. We first briefly revisit the reduced-order LTI model of the aggregate demand of a population of ACs in Section II and described the assumed population and control configuration for the present paper in Section III. In Section IV we present the three estimation approaches considered, and evaluate their accuracy and consistency during a series of LC events run on the proposed LC case study. Section V summarises the paper and provides concluding remarks.

II. Modelling dynamic aggregate demand response of a population of ACs

In this paper we consider the most common type of residential ACs: devices operating a "binary" compressor (either running at a fixed speed or not running). Such devices regulate temperature using a thermostat and a control mechanism known as *bang-bang control with hysteresis* or *on-off control*, which works as follows. The compressor is switched on when the thermostat reading exceeds a predefined value θ_+. With the compressor now engaged, the temperature starts to drop until the thermostat reading is below another predefined value θ_- (where $\theta_- < \theta_+$). At this point, the compressor turns off and the temperature gradually rises until it reaches θ_+ and the cycle starts again. This cyclic behaviour can be seen in Figure 1, which shows an example of temperature and power data from our AC hardware testing (in this case, a Delonghi DECON28AUP). Because the compressor is responsible for most of the energy usage of an AC unit, its cycling can be easily seen in the power curve. We also observe in the figure that when the compressor is on, the temperature inside the room gradually drops to $\theta_- \approx 17$ and when the compressor is off, the temperature rises to $\theta_+ \approx 17.9$.

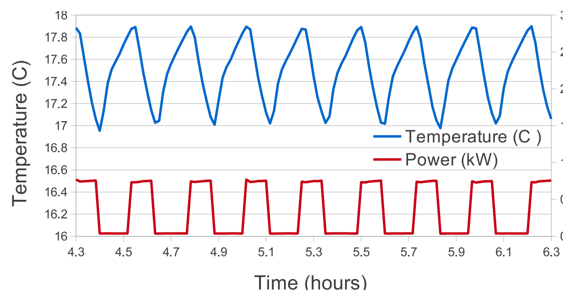

Fig. 1. Power and temperature over time of a Delonghi DECON28AUP air conditioning system tested at the CSIRO Newcastle testing facility. The temperature outside the test room was 27 ^{o}C.

Assuming a population of n such ACs, the dynamics of the i-th AC in the population (i.e., those shown in Figure 1) can be modelled by the well-known and extensively-used hybrid state model

$$\begin{aligned} \frac{d\theta_i(t)}{dt} &= -\frac{1}{C_i R_i}\big[\theta_i(t) - \theta_a(t) + m_i(t) R_i P_i - \alpha_i(t)\big], \\ m_i(t+\Delta t) &= \begin{cases} 0, & \text{if } \theta_i(t) \le \theta_i^- + u(t), \\ 1, & \text{if } \theta_i(t) \ge \theta_i^+ + u(t), \\ m_i(t), & \text{otherwise,} \end{cases} \end{aligned} \tag{1}$$

presented in [25], where Δt is an arbitrarily small time interval, $\theta_i(t)$ is the room temperature, θ_a is the ambient temperature outside the rooms (°C), C_i and R_i are the i-th room thermal capacitance (kWh/°C) and thermal resistance (°C/kW), and P_i is the cooling thermal power of the i-th AC (kW). The binary variable $m_i \in \{0,1\}$ ($i \in \{1,2,\ldots,n\}$,) represents the state of the compressor which switches *on* the AC ($m_i = 1$) or *off* ($m_i = 0$) to maintain the temperature θ_i within the pre-specified hysteresis band $[\theta_i^-, \theta_i^+]$, centred at $\theta_i^{\mathrm{r}} = (\theta_i^- + \theta_i^+)/2$. The noise process α_i represents thermal disturbances not modelled explicitly, and u is the proposed control signal (common to all ACs) to introduce small temporary temperature set-point offsets to the population during LC events.

The aggregate electrical power demand (kW) of the population of ACs is given by

$$d^{\mathrm{ac}}(t) = \sum_{i=1}^{n} m_i(t) \frac{P_i}{\mathrm{COP}_i}, \tag{2}$$

where COP_i is the coefficient of performance of the i-th AC, defined as the nominal ratio of rate of heat removal to electric power demand[1].

In [16], [23] we showed that, normalised to the maximum demand of the population, the aggregate demand response $d^{\mathrm{ac}}(t)$ to a common step change in temperature set points can be closely approximated by the LTI step response

$$d^{\mathrm{ac}}(t) \approx d^{\mathrm{ss}}(\theta^{\mathrm{r}}) - d^{\mathrm{g}}(t) \quad (3)$$

where $d^{\mathrm{ss}}(\theta^{\mathrm{r}})$ is the asymptotic steady state demand response of the population to constant temperature reference set points (with mean value θ^{r}), and $d^{\mathrm{g}}(t)$ is the unit step response of a second order transfer function model of the form

$$G(s) = \frac{b_2 s^2 + b_1 s + b_0}{s^2 + 2\xi\omega_n s + \omega_n^2}. \quad (4)$$

The model parameters b_2, b_1, b_0, ω_n and ξ in (4) are given as explicit functions of the means and variances of the physical parameters (such as thermal power, resistance and capacitance) distributed in the modelled population of ACs. For illustration (see [16], [23] for the full set of formulas),

$$\omega_n = \frac{\pi\mu_v}{(\sqrt{6}-a)\sqrt{1-\xi^2}}, \text{ and } \xi = \frac{\log(r)}{\sqrt{\pi^2 + \log^2(r)}}, \quad (5)$$

where $\mu_v = \mathrm{mean}\left[(\theta_a - \theta^{\mathrm{r}})/(CR)\right]$, $a = \exp\left[\log(\sqrt{2}) - \sigma_{\mathrm{rel}}^2 \frac{\log(3)}{\log(2)}\right]$, and

$$r = \left| \frac{\mathrm{erf}\left[(0.9 + \sqrt{8}\sigma_{\mathrm{rel}})^{-1}\right] - \frac{1}{2}}{\mathrm{erf}[0.9^{-1}] - \frac{1}{2}} \right|,$$

where $\mathrm{erf}[\cdot]$ is the Gauss error function.

These expressions for the parameters are obtained in [16], [23] by *analytically* fitting the step response of the second order transfer function (4) to a simplified parametric stochastic model for the aggregate demand of a distributed population of n ACs. Such demand can be calculated at each instant t by using (2) normalised to the maximum demand of the population.

In [16], [23], it was shown that the model (4) can closely approximate the aggregate power response of populations of ACs under a broad range of situations beyond the assumptions made to develop such model. For example, Figure 2 shows that the modelled aggregate power response obtained from (3) closely approximates the simulated responses $d^{\mathrm{ac}}(t)$ for populations comprising n =60 and n =10,000 ACs obtained from (2) using the distribution of parameters values in Table I. The input signal $u(t)$ in all cases is a temperature set point step increment of $0.5°$C at time $t = 0$.

Perhaps the main benefit of a second-order LTI model such as (4) is that it constitutes a very practical trade-off between fit and complexity: with only two states, it has the simplest structure capable of capturing the characteristic oscillations that appear in the aggregate demand of TCLs after a synchronisation event. This simplicity is a key enabler for the design of robust feedback control strategies for LC. In fact, we applied the LTI model (4) to successfully design an internal model controller in [16], [23] and an integral controller in [26], [23], both of which regulate the aggregate demand of a population of ACs. In the present paper, we show a different application for such a model: the estimation of energy savings during a LC event using a Kalman Filter. Note that the implementation of a Kalman Filter based on the reduced-order model (4) is substantially much simpler than using a hybrid model with potentially hundreds of states, as is the case when (2) is computed by replicating multiple instances of the hybrid model (1).

[1]For simplicity we assume a constant COP. However, in reality the COP varies mainly as a function of the temperature differential $\theta(t) - \theta_a$.

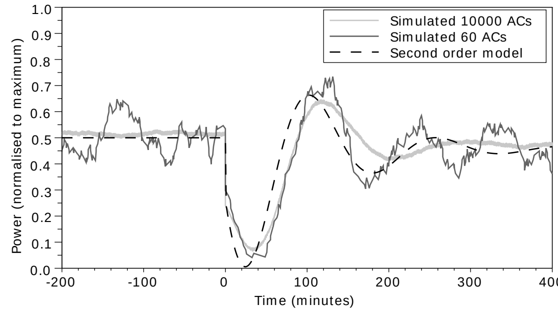

Fig. 2. Power response simulated from (1), (2) for small and large populations of ACs (60 and 10000 devices respectively) to a 0.5 C step change in temperature set point, and modelled LTI response simulated from (3), (4) with parameters computed with the formulas given in [16], [23].

TABLE I
SIMULATION PARAMETERS.

Param.	Value	Description
R	2	Mean thermal resistance (log-normal distribution with normalised variance σ_{rel}^2) (°C/kW)
C	3.6	Mean thermal capacitance (log-normal distribution with normalised variance σ_{rel}^2) (kWh/°C)
P	6	Mean thermal power (log-normal distribution with normalised variance σ_{rel}^2) (kW)
θ^{r}	20	Mean temperature set point for the ACs uniformly distributed in $[19.5, 20.5]$ (°C). $\theta_i^{\mathrm{r}} = (\theta_i^- + \theta_i^+)/2$.
H	1	Hysteresis width $(\theta_i^+ - \theta_i^-)$ (°C)
θ_a	Variable	Ambient temperature (°C)
σ_w^2	0.01	Variance of the noise process w in (1)
σ_{rel}	0.2	Standard deviation of log-normal distributions as fraction of the mean value for R, C and P
COP	2.5	Coefficient of performance (thermal on electrical power)
n	60	Number of ACs in the population

III. CASE STUDY

During a LC event, the aggregate power demand of the ACs in the population may not be easily available due to communication costs and privacy issues. On the other hand, what is readily available is the aggregate demand $d(t)$ of the users connected at the electrical distribution feeder that powers the ACs. Thus we can express this demand as $d(t) = d^{\mathrm{ac}}(t) + W(t)$ where $d^{\mathrm{ac}}(t)$ is AC controllable load and $W(t)$ is the rest of the loads (uncontrollable). For the case study presented in this paper, uncontrolled loads $W(t)$ correspond to real demand data logged from a distribution feeder located in a suburban area on the east coast of Australia. The ACs are simulated in PowerDEVS [27] using real ambient temperature data $\theta_a(t)$ logged at a nearby location in the same period. These power and temperature readings are the actual

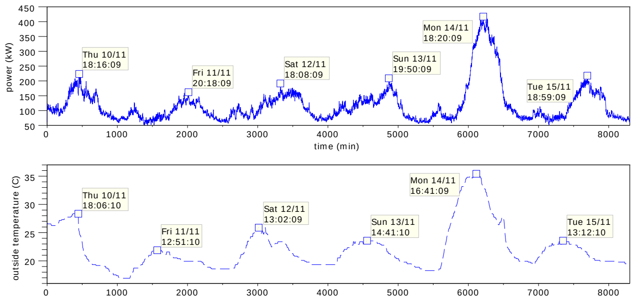

Fig. 3. Top: power demand of 70 residential customers in a newly-developed residential area in Australia during six days in November 2011. Bottom: ambient temperature. Time of daily peaks is displayed.

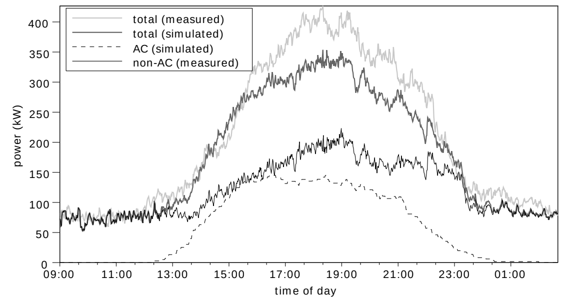

Fig. 4. Real and simulated power load of 70 residential customers on a hot day with a maximum temperature of 35.5 C.

data corresponding to six consecutive days at a residential distribution feeder supplying 70 houses, and are shown in Figure 3.

We simulate the ACs using (2) and the parameters in Table I, which were adapted from [15] to the characteristics of Australian suburban houses. They were obtained as follows. The mean thermal capacitance C was calculated as indicated in [15] assuming a 90 m^2 building. The mean thermal resistance R was calculated assuming 90 m^2 of insulated tiled roof, 90 m^2 of timber floor with carpet, 100 m^2 of brick walls and 14 m^2 of single glass windows (the heat transfer coefficients used for these materials were taken from the AIRAH Handbook [28]). The mean thermal power P was calculated using the AIRAH "FairAir" Calculator [29], assuming proximity to Sydney, 3 occupants and the same parameters used to calculate thermal resistance (considering 3 m ceiling height and 3, 3, 4 and 4 m^2 of internally-shaded windows facing East, West, North and South respectively). To simulate different types of ACs and user-defined temperature set points, θ^- and θ^+ are assumed uniformly distributed in the population. The mean hysteresis width (i.e., $\theta^+ - \theta^-$) considered is 1 ^oC. The rest of the parameters were obtained from [15].

Note the strong correlation between temperature and demand in Figure 3: the load peak of Monday (maximum temperature 35.5 C) was practically twice as much as that of the next day (maximum temperature 23.5 C).

Mild-temperature week days such as Friday 11/11 and Tuesday 15/11 in Figure 3 are assumed without AC load. We then construct the load profile for the hot Monday 14/11 by assuming that the temperature uncorrelated loads are equal to the total load of a mild-temperature day, and aggregating the (PowerDEVS simulated) AC load using the temperature profile for Monday 14/11.

Figure 4 shows the simulated load using the temperature profile of the hot Monday 14/11 as well as the actual load of that day. The temperature-uncorrelated load component (non-AC) for the simulation was assumed to be the total load of the mild Tuesday 15/11. We assume that 60 out of the 70 customers in the residential area have an AC operating during the simulated day. This assumption aligns with studies that indicate that 90% of the houses in the Sydney area operate their ACs on the warmest days [30]. The ACs are randomly turned on from 12:30 to 16:00 to simulate the arrival of people at their houses, and turned off between 21:00 and 00:30, as people go to sleep. These hours align with times reported in a previous study on residential AC usage in New South Wales, Australia [31].

Note that the simulated and measured loads in Figure 4 share the following fundamental characteristics: the peak demand occurs at the same time of the day and the rate of load increase in the afternoon and decrease in the evening are very similar. The difference between the simulated and measured loads accounts to approximately 1 kW per customer during the peak hours. This difference may be attributed to other temperature-dependant loads that were not simulated (such as fridges and freezers, ceiling and floor fans) and the discrepancy between the assumed representative parameters in Table I[2]. and the real population.

Because we do not have access to the actual physical parameters of these houses (nor do we have the ability to control their ACs at this point), in the remainder of this paper we will use the simulated load in Figure 4 to compare against the case when a LC event takes place. This has the advantage of allowing us to know exactly "what would have happened" had the LC not occurred, which is otherwise impossible to find out in a real scenario.

Let us now revisit the feedback LC strategy considered in [26]. The LC set up under consideration is illustrated in Figure 5, where a small number of houses with controllable temperature set-point ACs are connected to a common power distribution point. The total aggregate demand of the houses (including that of ACs and other, uncontrolled loads) is measured at the distribution supply point and transmitted to a central controller, which uses it to compute new temperature set-point offsets broadcast to the ACs. The controller-ACs communications are one-way, as the ACs do not need to transmit their demand.

Figure 6 presents a block diagram of this feedback scheme. The controller uses the error between the desired aggregate demand $d^{\mathrm{r}}(t)$ for the feeder during the LC event, and the measured demand $d(t)$, which aggregates the demand of the ACs, $d^{\mathrm{ac}}(t)$, and that of non-controllable loads, $W(t)$. The computed temperature set point offset $u(t)$ is broadcast as a common control signal to the population of ACs.

In the following section we use the LTI model (4) to estimate energy savings during LC events. We implement the LC events using the integral controller proposed in [26]. This

[2] In a real implementation, these parameters could be estimated from the actual population by performing system identification experiments, as done in [15] and suggested in [26].

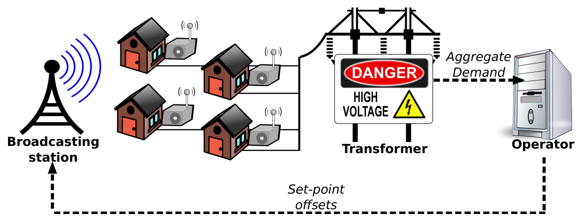

Fig. 5. One-way communications LC set up. The houses with controllable ACs are connected to a common distribution feeder. The total aggregated power demand of the houses is measured at the distribution supply point and fed back to the controller, which computes and broadcasts new set-point offsets to the ACs.

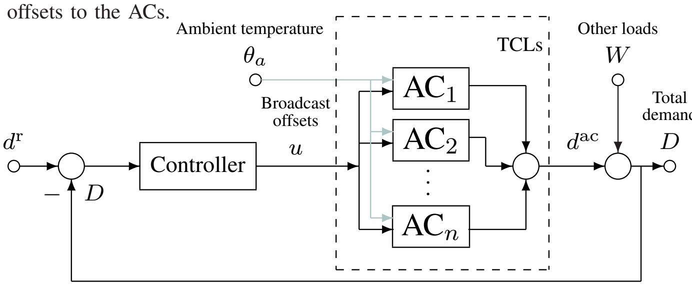

Fig. 6. Schematics of a feeder load feedback control of a population of ACs. Based on the tracking error $d^{\mathrm{r}}(t) - D(t)$, the controller computes the temperature set point offset $u(t)$ broadcast to the ACs.

controller successfully regulates the aggregate demand of a population of ACs in the configuration shown in Figures 5 and 6. Note, however, that our technique for savings estimation during a LC event is completely independent from the control strategy used to implement the scenario. Thus, the approach in this paper could also be applied to other control techniques such as [15], [13], [32].

IV. Model-based estimation of energy reductions in a LC event.

Accurate estimates of how much energy is effectively saved during LC events provide valuable information to assess the financial benefits of LC [4]. Note that such savings can only be estimated, as one would otherwise require to know what the demand would have been had the LC event not taken place.

If a forecast of the uncontrolled demand is available from the utility (which may be estimated from historical data, a regression analysis or usage statistics of a control group) LC savings can be estimated as the difference between the forecast and the real power demand. However, at feeder-level scenarios such as the one described in the previous section, such forecast may not be available, or may not be accurate enough.

In this section we compare estimates of energy savings during LC events using three methods that do not require a load forecast. In particular, one of these methods (the main contribution of this paper) uses a Kalman filter that incorporates the model (4) to generate real-time predictions. We describe these methods in more detail next and then compare their performance on a LC event built using the case study presented in the previous section.

Note that the discussed methods represent causal solutions that only take into account the information available until present time to compute the estimates. While one could use non-causal methods that take into account the power demand before *and after* the LC event to calculate the energy savings, causal methods provide the advantage of estimating savings in real time. Real-time savings estimates are very desirable in situations when there is a feedback strategy in place that computes the control signal based on the achieved savings until present. In other words, one could use causal methods (as the one proposed in this paper) to develop controllers to manipulate power in real time to achieve a certain level of desired energy savings.

A. Energy savings estimation methods

Let d_k denote the sampled total power demand (in kW), averaged over the k-th one-minute interval, read at the feeder to which the ACs are connected. Similarly, d_k^{ac} and W_k denote sampled average values over the k-th one-minute intervals. Three energy saving estimation methods are considered:

1) Mean: Aggregate demand before the LC event (referred to as *non-LC demand*) is assumed equal to the average of the last 30 minutes before the LC event, namely $\bar{d} = \sum_{j=t_s-30}^{t_s-1} d_j/30$. Energy savings between times t_1 and t_2 are then estimated as

$$\hat{S} = \sum_{k=t_1}^{t_2} \left(\bar{d} - d_k\right)/60. \tag{6}$$

This estimation method is mentioned in [22] as a candidate for periods of “flat” load, but regarded as inadequate for periods of increasing of decreasing loads.

2) Linear regression (LR): Aggregate non-LC demand is computed as a linear regression of the last 60 readings before the event, $[d_{t_s-60}, d_{t_s-59}, \ldots, d_{t_s-1}]$. Energy savings between times t_1 and t_2 are then estimated as

$$\hat{S} = \sum_{k=t_1}^{t_2} (ak + b - d_k)/60, \tag{7}$$

where a, b are the fit parameters for the linear regression. This method was used to estimate the energy savings in several real-world LC trials [22].

3) Kalman filter (KF): The KF method uses a Kalman Filter to combine *predicted* information about the population of ACs (this information comes in the form of model structure) with *observed* information about such population (i.e., measurement data). In particular, our method uses the results obtained by exciting the model (3)-(4) with u_k as the predicted information; and aggregate power demand measurements d_k as observed information. Using a Kalman Filter to estimate energy savings in a LC scenario is the main contribution of this paper.

Recall [33] that the Kalman filter addresses the problem of estimating the state x of a controlled discrete-time process with l inputs, p outputs and q states, assuming this process is governed by the linear stochastic state-space difference equation

$$x_k = \mathbf{A}x_{k-1} + \mathbf{B}u_k + w_{k-1} \tag{8}$$

with a measurement equation

$$d_k = \mathbf{C}x_k + \mathbf{D}u_k + v_k \tag{9}$$

where $x_k \in \mathbb{R}^{q\times 1}$, $d_k \in \mathbb{R}^{p\times 1}$, $u_k \in \mathbb{R}^{l\times 1}$, $\mathbf{A} \in \mathbb{R}^{q\times q}$, $\mathbf{B} \in \mathbb{R}^{q\times l}$, $\mathbf{C} \in \mathbb{R}^{p\times q}$ and $\mathbf{D} \in \mathbb{R}^{1\times l}$. The random variables $v_k \in \mathbb{R}^{p\times 1}$ and $w_k \in \mathbb{R}^{q\times 1}$ are measurement and process noises assumed to be zero-mean and normally distributed with covariance matrices V and Q.

We refine our model of the total demand d_k as

$$d_k = d_k^{\mathrm{ac}} + W_k^p + W_k^c, \tag{10}$$

where d_k^{ac} is the variation in demand of the controlled loads (ACs) as a direct consequence of the control action (that is, if the control signal $u_k = 0\ \forall t$, then $d_k^{\mathrm{ac}} = 0$), and W_k^p represents the daily oscillation in power demand for both temperature related and independent loads minus its mean, while W_k^c represents a baseline load constant over the whole period of time.

The representation of demand in (10) is embedded in the KF scheme for state estimation in the state-space representation by rewriting (8) and (9) as

$$\begin{bmatrix} x_k^a \\ x_k^p \\ x_k^c \end{bmatrix} = \begin{bmatrix} \mathbf{A}_a & 0 & 0 \\ 0 & \mathbf{A}_p & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{k-1}^a \\ x_{k-1}^p \\ x_{k-1}^c \end{bmatrix} + \begin{bmatrix} \mathbf{B}_a \\ 0 \\ 0 \end{bmatrix} u_k + w_k, \tag{11}$$

$$d_k = \begin{bmatrix} \mathbf{C}_a & \mathbf{C}_p & 1 \end{bmatrix} \begin{bmatrix} x_k^a \\ x_k^p \\ x_k^c \end{bmatrix} + \mathbf{D}_a u_k + v_k. \tag{12}$$

where x_k^a, x_k^p and x_k^c are vectors representing the internal state of the system during the kth time interval. The vector x_k^a describes the states associated with controlling the ACs during a LC event, the vector x_k^p represent the states associated with the daily oscillation commonly observed in power consumption, the single-element vector x_k^c represents the state associated with the constant, base-load demand, and the vector w_k represents the process noise. Let us describe the space-state subsystems associated with each of these state vectors.

The subsystem

$$x_k^a = \mathbf{A}_a x_{k-1}^a + \mathbf{B}_a u_k + w_k^a, \quad \text{with} \quad d_k^{\mathrm{ac}} = \mathbf{C}_a x_k^a + \mathbf{D}_a u_k + v_k^a, \tag{13}$$

is the (1 minute sampled) discretised state-space representation of the LTI model (4), which describes the power demand of the population of ACs controlled by set point offsets u_k.

The zero-mean daily power oscillation W^p is described by

$$x_k^p = \mathbf{A}_p x_{k-1}^p + w_k^p, \quad \text{and } W_k^p = \mathbf{C}_p x_k^p + v_k^p.$$

We model W^p as a sinusoidal wave with frequency $\mu = 1.16 \times 10^{-5}$ Hz (once a day), which represents the daily variability in demand. This parameter choice is evident from Figure 7, which shows the power spectral density of the demand data in Figure 3 (top). The peak seen in Figure 7 confirms the intuition that the most important frequency in the aggregate power demand is the daily oscillation caused by people's routine: every week day, the demand rises towards the evening, when most people arrive home, turn on the T.V., AC, start cooking, etc. Conversely, the minimum power demand occurs during the early hours of the morning, when most people are asleep. Note that this daily oscillation has a small frequency variation over time (evidenced by the broad base of the peak in the power density plot) but it can be considered as the dominant frequency for a window of a few consecutive days.

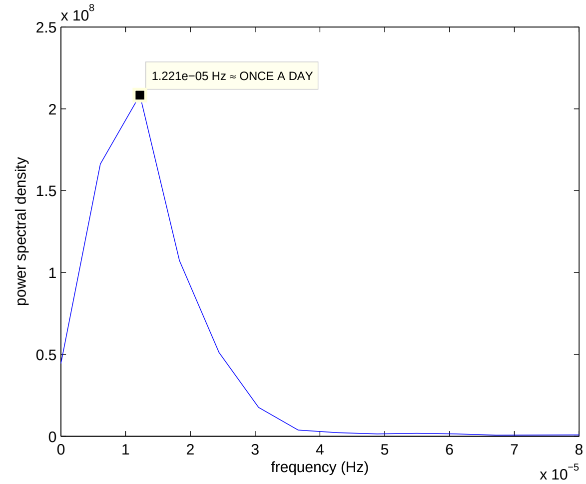

Fig. 7. Power Spectral Density of the power signal in Figure 3 generated using the `pwelch` MATLAB® command.

The discrete state-space matrices sampled at 1-minute intervals of this sinusoidal wave W^p with frequency $\mu = 1.16 \times 10^{-5}$ are

$$\mathbf{A}_p = \begin{bmatrix} \cos(2\pi 60\mu) & \sin(2\pi 60\mu) \\ -\sin(2\pi 60\mu) & \cos(2\pi 60\mu) \end{bmatrix}, \mathbf{C}_p = \begin{bmatrix} 1 & 0 \end{bmatrix}.$$

Lastly, the constant baseline W^c can be modelled as

$$x_k^c = x_{k-1}^c + w_k^c, \quad \text{and } W_k^c = x_k^c + v_k^c.$$

By combining available power measurements, the model of the system and stochastic information about the noises perturbing the process and the measurements, the Kalman filter provides the best linear (optimal) estimate $\hat{x}$ of the state x [33]. This estimate is calculated as

$$\hat{x}_k = \hat{x}_k^- + \mathbf{K}_k(d_k - \mathbf{C}\hat{x}_k^- - \mathbf{D}u_k) \tag{14}$$

where

$$\hat{x}_k^- = \mathbf{A}\hat{x}_{k-1} - \mathbf{B}u_k, \tag{15}$$

and $\mathbf{K}_k$ is the Kalman filter gain or *blending factor*, which is recursively computed to minimise the covariance of the estimation error by solving a standard Riccati difference equation (see e.g., [33]).

We estimate the output d_k^{ac} associated with the ACs using

$$\widehat{d_k^{\mathrm{ac}}} = \mathbf{C}_a \hat{x_k^a} + \mathbf{D}_a u_k, \tag{16}$$

where $\hat{x_k^a}$ is the state subset of $\hat{x}_k$ associated with the effect of controlling the ACs (i.e., an estimate of x_k^a in (11)).

From (16), the energy savings between times t_1 and t_2 can be computed as

$$\hat{S} = 60 \sum_{i=t_1}^{t_2} \widehat{d_i^{\mathrm{ac}}}. \tag{17}$$

Extensive numerical tests indicate that the reduction estimate $\hat{S}$ from (17) varies greatly depending on the chosen process and measurement error covariance matrices Q and V for the process and measurement noises w_k and v_k in (8), (9) (these matrices are used to compute $\mathbf{K}_k$ in (14)). Thus, use

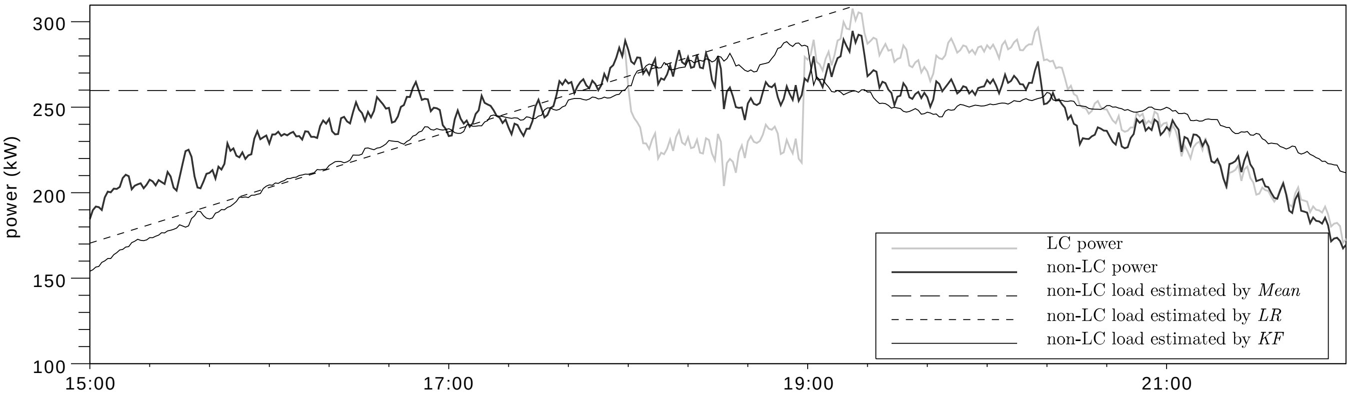

Fig. 8. Comparison of estimated energy savings by the proposed methods for the LC scenario from 18:00 to 19:00. Dashed straight lines: assumed non-LC consumption used in the Mean and LR methods. Continuous thin black line: assumed non-LC load for the KF method.

a of KF approach requires careful tuning of these matrices. For the results presented in the following section, we have manually adjusted these matrices to

$$Q = \begin{bmatrix} 3025 & 1925 & 0 & 0 & 0 \\ 1925 & 1225 & 0 & 0 & 0 \\ 0 & 0 & 25 & 15 & 0 \\ 0 & 0 & 15 & 9 & 0 \\ 0 & 0 & 0 & 0 & 0.99 \end{bmatrix} \quad \text{and} \quad V = 0.064.$$

While finding optimal weighting matrices for Q and V falls beyond the scope of the present paper, techniques such as the auto-covariance least-squares method may be used for this purpose [34].

The alternative estimates of energy savings given in (6), (7) and (17) are interpreted graphically in Figure 8. The dashed straight lines are the assumed non-LC consumption used in (6) and (7). The continuous thin black line is the assumed non-LC load used for the KF estimate (17), which is the only one that effectively captures the daily oscillation of the non-LC load. The estimated savings are computed by calculating the area between the assumed non-LC load of the method and the LC load (solid grey). These estimates can be compared with the actual savings, which are represented by the area between the actual non-LC load (non measurable in reality, but available in a simulated scenario), and the LC load.

We next compare the performance of these estimation methods on a series of simple LC events.

B. Case study: estimation performance comparison

Four LC scenarios were tested on the data from the case study presented in Section III, in which the aggregate power (ACs and other loads) measured at the feeder is controlled with the integral controller presented in [26] ($K_I = 2$). In all of these scenarios, the aim of the controller is to reduce the total power by a 20% of the demand at the start of the event, and maintain the reduction for a given period of time. The four scenarios considered differ in their duration, maintaining the target demand reduction for 1, 2, 3 and 4 hours, at different times of the day. For each type of scenario, we simulate a series of (independent) LC events. In each simulation, we assume that no other LC event took place during that day and that all 60 ACs are controllable. The first LC event for each scenario starts at $t_s =$14:00, the second one at $t_s =$14:30 and the last one finishes at $t_e =$22:00 (e.g., the 3-hour events are: 14:00-17:00, 14:30-17:30, . . . , 19:00-22:00).

Table II shows the real and estimated savings for each simulated event for the 1, 2, 3 and 4-hour scenarios respectively. The simulations were run with the same population parameters and non-AC loads as Figure 4. Each table shows the real savings for the event and the estimates $\hat{S}$ computed for each method as described in Section IV-A above between $t_1 = t_s$ and $t_2 = t_e + 60$ min (savings during the reduction time and "negative" savings for the following hour[3]). The underlined estimates are the closest to the real savings.

It can be seen in Table II that the KF estimation, which incorporates the proposed LTI model (4), yields the closest estimates to the real savings in most cases. Furthermore, Figure 9 displays, for each estimation method, the number of *satisfactory* estimates, where we classify an estimate as satisfactory when it is within plus or minus 20 percent of the real savings S (i.e., $0.8\hat{S} \leq S \leq 1.2\hat{S}$). The histograms show that for all four reduction periods considered (1, 2, 3 and 4 hours), the KF estimation achieves the largest number of satisfactory estimates. Note that the longer the reduction period considered, the larger the number of satisfactory estimates, which is expected, as the KF estimation is the only method that takes into account daily load variation.

The performance of the KF estimates is further illustrated in Figure 10, which depicts the simulated scenario corresponding to the 2-hour reduction LC event on Table II from 17:00 to 19:00. Such scenario, where the AC load is reduced to roughly half of the maximum AC load during two hours is of practical relevance, and aligns with recent LC trials run in Perth (Australia) [5].

In the top plot of Figure 10 we see that the non-LC load between 17:00 and 19:00 does not continue in a straight line, either as a constant or with a slope, which explains why Mean and LR are particularly inaccurate for this event.

[3] Note that one may be interested in estimating only the savings during the event (i.e., for peak shifting applications) whether others might prefer to account for several hours after the event (i.e., to quantify net savings). Although we chose 60 minutes as a middle point for these two scenarios, the user could otherwise easily change this value as needed for the application.

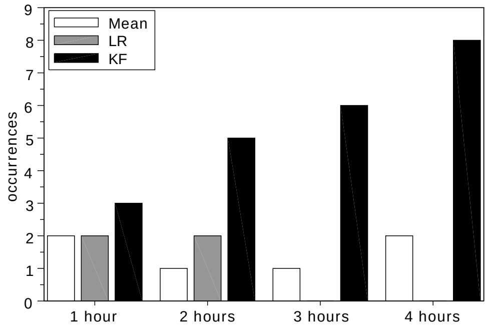

Fig. 9. Number of times each estimation methods is within plus or minus 20 percent of the real savings S as per Table II.

For comparison, the bottom plot shows the breakdown of states estimates in the KF model (12), $\hat{x}_k^c$ (baseline), $\mathbf{C}_p\hat{x}_k^p$ (daily oscillations), and $\mathbf{C}_a\hat{x}_k^a+\mathbf{D}_a u_k$ (ACs), together with an additional estimate (ACs model) based simply by driving the model (4) with the control signal u_k. Note that while the latter estimate may be acceptable during the LC period, it produces a large error peak once the LC period ends (a phenomenon similar to the "cold load" pickup [35], [36]). This large peak is due to the fact that the second order model (4), since it is linear, does not incorporate the saturation in aggregate power when all of the ACs are turned on (or off). Note that if the model was perfect, we could just estimate the savings of a LC event by exciting such model with the control signal sent to the ACs and compute the difference between that and zero. However, this open-loop approach is of a much more brittle nature than a feedback estimation approach implemented by the Kalman Filter, where measurements are used to correct the estimates in real-time. In contrast, the KF compensates for the daily dynamics $\hat{W}^p$ (as shown in Figure 8), by optimally balancing in the estimates the information from the model and that from real measurements.

Note that the proposed KF approach divides the load in three categories (baseline, oscillations and ACs) even prior to the control event. Before the event, the control signal u_k is equal to zero and the expected change in the ACs due to LC is also zero. We can see in the bottom plot of Figure 10 that the value for the estimated ACs only starts changing significantly after the beginning of the event (i.e., 17:00). Moreover, starting the filter beforehand helps "warm up" the KF, so that when the event starts, initial transients minimised and the baseline and daily oscillations estimates are clearly separated.

The accuracy of our KF estimate depends on two main factors: how well the LTI model (4) represents the aggregate power of the population and how observable the effect of LC is wherever the power measurements are taken (in the case of this paper, at the feeder). Regarding how closely the LTI model (4) represents the power demand of the population, it was shown in [16] that the dynamics are successfully captured for an operational range well beyond the one assumed to develop the model. Nevertheless, if a more accurate (and presumably more complex) model was required, the approach presented in the present paper could still be applied by using the desired model instead of the subsystem (13). In terms of the observability of the AC dynamics at the feeder level, the larger the proportion of controlled ACs in the total load, the more observable the changes caused by LC will be. Fortunately, LC is more likely to take place during high-temperature days, when ACs represent a significant part of the total load.

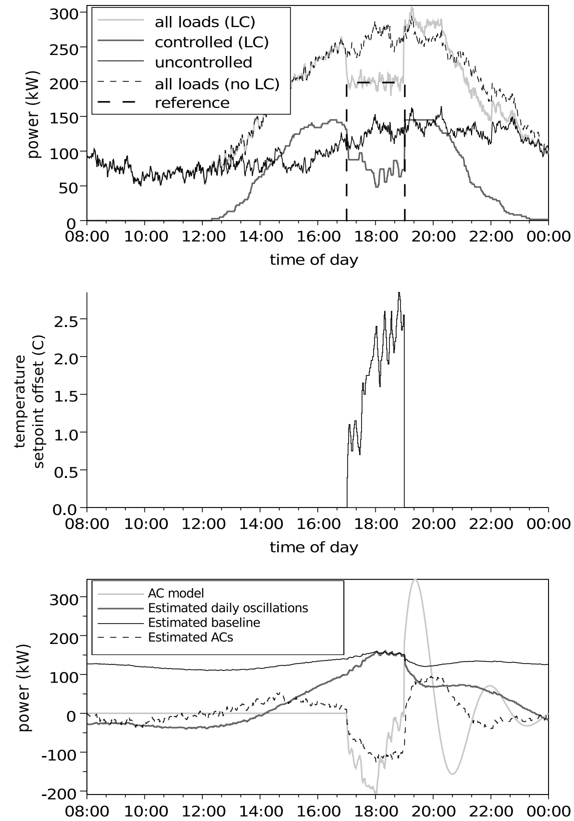

Fig. 10. LC event (from 17:00 to 19:00) where KF returns the best estimates. Top: Power with and without LC, and breakdown of LC-power into controlled and uncontrolled loads. Middle: Control signal calculated by the integral controller. Bottom: Second-order LTI model response and KF breakdown of LC power into baseline, daily oscillation and AC variation due to LC.

V. Conclusions

It is important to accurately estimate energy savings associated with LC in peak demand reduction or load tracking events. We have proposed an estimation strategy based on a Kalman filter that integrates the LTI model developed in [16]. We have evaluated the accuracy of the proposed estimation approach on a series of LC events constructed on the case study considered. The estimates based on the Kalman filter are consistently better than simple estimates based on mean demand and linear regressions, and only require distribution feeder readings, which may aggregate uncontrolled demand. Importantly, unlike most estimation methods in the literature, the Kalman filter estimate does not require extensive historical data or a control group.

In our view, the results presented in this paper highlight the practical relevance of a reduced-order model such as the one in [23] in the application of powerful and well-understood control and estimation methods for load management of TCLs.

References

[1] K. Schisler, T. Sick, and K. Brief, "The role of demand response in ancillary services markets," in *IEEE/PES Transmission and Distribution Conf. and Exposition. IEEE/PES*, April 2008.

[2] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, January 2011.

TABLE II
REAL AND ESTIMATED SAVINGS (KWH) FOR THE CONSIDERED LC EVENTS. BEST ESTIMATE UNDERLINED.

Start time	1-hour LC event				2-hour LC event				3-hour LC event				4-hour LC event			
	Real	Mean	LR	KF	Real	Mean	LR	KF	Real	Mean	LR	KF	Real	Mean	LR	KF
14:00	-56.2	69.8	**-57.1**	-14.0	-153.6	89.3	**-175.1**	-90.8	-277.1	90.8	-360.6	**-209.8**	-411.8	97.0	-590.7	**-393.4**
14:30	-55.7	54.4	-112.6	**-5.5**	-144.0	58.3	-290.5	**-76.7**	-256.0	59.6	-536.4	**-171.8**	-383.2	44.0	-864.7	**-319.9**
15:00	-60.0	20.6	-118.5	**-14.8**	-150.0	-3.6	-297.9	**-100.7**	-256.2	-28.2	-534.6	**-198.5**	-369.3	-60.4	-835.9	**-303.5**
15:30	-50.6	21.0	**-51.4**	-23.7	-127.9	10.9	**-141.6**	-95.0	-222.1	-17.8	-279.7	**-187.0**	-307.6	-43.3	-443.7	**-260.5**
16:00	-47.2	12.6	-80.4	**-36.0**	-113.7	-4.0	-198.7	**-104.0**	-186.8	-27.9	-361.1	**-172.0**	-251.8	-61.5	-570.0	**-230.0**
16:30	-44.8	-2.1	-89.5	**-46.6**	-114.7	-41.8	-226.3	**-120.0**	-174.9	-77.3	-394.4	**-178.5**	-230.5	-137.7	-623.0	**-225.9**
17:00	-42.6	-25.5	-89.5	**-59.4**	-101.9	-68.8	-204.4	**-121.2**	-157.8	-126.7	-360.1	**-174.6**	-190.6	-199.1	-556.9	**-190.7**
17:30	-64.3	-15.3	8.4	**-90.6**	-127.4	-60.3	-8.6	**-158.3**	-186.7	-131.0	-40.7	**-205.2**	-208.0	**-205.6**	-66.1	-193.7
18:00	-24.2	-12.1	-96.3	**-35.1**	-67.0	**-66.9**	-244.6	-87.1	-122.7	-172.5	-477.7	**-112.7**	-106.3	-238.9	-705.9	**-58.3**
18:30	-35.9	-60.7	-132.4	**-46.8**	-90.3	-159.7	-311.1	**-96.9**	-121.0	-276.9	-537.2	**-93.8**				
19:00	-60.6	**-62.0**	6.4	-83.5	-122.6	-171.5	-25.7	**-120.2**	-140.2	-269.6	-17.8	**-90.2**				
19:30	-80.7	-129.0	-164.2	**-83.7**	-108.3	-237.3	-313.5	**-80.2**								
20:00	-79.1	-105.1	-4.2	**-67.2**	-68.4	-159.3	56.4	-21.0								
20:30	-72.2	-115.0	**-53.8**	-44.0												
21:00	-45.4	**-54.2**	83.1	0.3												

[3] M. Bozchalui, S. Hashmi, H. Hassen, C. Canizares, and K. Bhattacharya, "Optimal operation of residential energy hubs in smart grids," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1755–1766, 2012.

[4] C. Dunstan, N. Ghiotto, and K. Ross, "Report of the 2010 survey of electricity network demand management in Australia," Alliance to Save Energy, Tech. Rep., 2011.

[5] A. Blaver, "Perth solar city annual report," Perth Solar City, Tech. Rep., 2011. [Online]. Available: http://www.perthsolarcity.com.au/annual-report/

[6] ETSA utilities. (2010) Demand management program interim report no. 3. [Online]. Available: www.etsautilities.com.au/public/download.jsp?id=11891

[7] G. Heffner, C. Goldman, B. Kirby, and M. Kintner-Meyer, "Loads providing ancillary services : Review of international experience," *Lawrence Berkeley National Lab.*, 2007.

[8] B. Kirby, J. Kueck, T. Laughner, and K. Morris, "Spinning reserve from hotel load response: initial progress," Oak Ridge Nat. Lab. and Tennessee Valley Authority, Tech. Rep., 2008.

[9] J. Eto, J. Nelson-Hoffman, E. Parker, C. Bernier, P. Young, D. Sheehan, J. Kueck, and B. Kirby, "Demand response spinning reserve demonstration—phase 2 findings from the summer of 2008," LNBL, Tech. Rep. LBNL-2490E, 2009.

[10] F. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proc. Power Systems Computation Conference*, 2011.

[11] 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 Conf.*, 2011.

[12] K. Kalsi, F. Chassin, and D. Chassin, "Aggregated modeling of thermostatic loads in demand response: A systems and control perspective," in *IEEE Conf. on Decision and Control*, 2011.

[13] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, "Modelling and control of thermostatically controlled loads," in *Proc. Power Systems Computation Conference*, Stockholm, 2011.

[14] B. Kirby and E. Hirst, "Load as a resource in providing ancillary services," *Proc. of the American Power Conf.*, 1999.

[15] 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*, vol. 50, no. 5, 2009.

[16] 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*, vol. 55, pp. 36–48, 2012.

[17] G. R. Newsham, B. J. Birt, and I. H. Rowlands, "A comparison of four methods to evaluate the effect of a utility residential air-conditioner load control program on peak electricity use," *Energy Policy*, vol. 39, no. 10, 2011.

[18] J. Bode, M. J. Sullivan, and J. H. Eto, "Measuring short-term air conditioner demand reductions for operations and settlement," LBNL, Berkeley, Tech. Rep., 01/2012 2012.

[19] C. Gellings, "Estimating the costs and benefits of the smart grid: A preliminary estimate of the investment requirements and the resultant benefits of a fully functioning smart grid," *Electric Power Research Institute (EPRI), Technical Report (1022519)*, 2011.

[20] I. Khanna, "Smart grid application: Peak demand management trial - the western australian experience," in *Innovative Smart Grid Technologies Asia (ISGT), 2011 IEEE PES*, 2011, pp. 1–6.

[21] M. E. Poulad, A. S. Fung, F. Rad, and T. Fotinakopoulos, "Potential energy saving due to peaksaving/lca process using ann in summer in toronto," in *EIC Climate Change Technology Conference*, 2013.

[22] J. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, and D. S. Watson, "Demand response spinning reserve demonstration," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-62761, 2007.

[23] C. Perfumo, "Dynamic modelling and control of heterogeneous populations of thermostatically controlled loads," Ph.D. dissertation, Electrical Engineering and Computer Science, The University of Newcastle, Callaghan, NSW, Australia, 2013.

[24] J. Mathieu and D. Callaway, "State estimation and control of heterogeneous thermostatically controlled loads for load following," in *System Science (HICSS), 2012 45th Hawaii International Conf. on*, jan. 2012.

[25] C.-Y. Chong and A. S. Debs, "Statistical synthesis of power system functional load models," in *Proceedings of the 18th IEEE Conference on Decision and Control*, vol. 18, December 1979, pp. 264–269.

[26] J. H. Braslavsky, C. Perfumo, and J. K. Ward., "Model-based feedback control of distributed air-conditioning loads for fast demand-side ancillary services." *Accepted for publication in the 52nd IEEE Conference on Decision and Control, Florence, Italy.*, 2013.

[27] F. Bergero and E. Kofman, "PowerDEVS. A Tool for Hybrid System Modeling and Real Time Simulation," *Simulation: Transactions of the Society for Modeling and Simulation Intl.*, vol. 87, 2011.

[28] Australian Institute of Refrigeration, Air-Conditioning and Heating (AIRAH), *Handbook – Millennium Edition*, 3rd ed. AIRAH, 2000.

[29] Australian Institute of Refrigeration, Air Conditioning and Heating Inc. (AIRAH). (2012) FairAir Calculator. http://www.fairair.com.au/Calculator.Size.aspx.

[30] M. Hart and R. de Dear, "Weather sensitivity in household appliance energy end-use," *Energy and Buildings*, vol. 36, no. 2, 2004.

[31] Energy Efficient Strategies. (2005) Status of air conditioners in Australia - updated with 2005 data. [Online]. Available: http://www.energyrating.gov.au/wp-content/uploads/Energy_Rating_Documents/Library/Cooling/Air_Conditioners/200509-ac-aust.pdf

[32] N. Lu and S. Katipamula, "Control strategies of thermostatically controlled appliances in a competitive electricity market," in *IEEE Power Engineering Society General Meeting*, vol. 1, June 2005, pp. 202–207.

[33] G. Welch and G. Bishop. An Introduction to the Kalman Filter. [Online]. Available: http://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf

[34] B. M. Akesson, J. B. Jorgensen, N. K. Poulsen, and S. B. Jorgensen, "A generalized autocovariance least-squares method for kalman filter tuning," *Journal of Process Control*, vol. 18, pp. 769–779, 2008.

[35] C.-Y. Chong and R. P. Malhamé, "Statistical synthesis of physically based load models with applications to cold load pickup," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-103, no. 7, 1984.

[36] S. Ihara and F. Schweppe, "Physically based modeling of cold load pickup," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-100, no. 9, pp. 4142 –4150, September 1981.