

Load management: Model-based control of aggregate power for populations of thermostatically controlled loads

Cristian Perfumo^{a,c,*}, Ernesto Kofman^b, Julio H. Braslavsky^a, John K. Ward^a

^a CSIRO Energy Technology, 10 Murray Dwyer Circuit, Mayfield West, NSW 2304, Australia

^b CIFASIS-CONICET, Department of Control, FCEIA-UNR, Riobamba 245 bis, 2000 Rosario, Argentina

^c Faculty of Engineering and Built Environment, The University of Newcastle, NSW, Australia

ARTICLE INFO

Article history:

Received 29 June 2011

Accepted 20 October 2011

Available online 22 November 2011

Keywords:

Demand side management

Load management

Air conditioning

Internal model control

Load aggregation

Control signal quantization

ABSTRACT

Large groups of electrical loads can be controlled as a single entity to reduce their aggregate power demand in the electricity network. This approach, known as load management (LM) or demand response, offers an alternative to the traditional paradigm in the electricity market, where matching supply and demand is achieved solely by regulating how much generation is dispatched. Thermostatically controlled loads (TCLs), such as air conditioners (ACs) and fridges, are particularly suitable for LM, which can be implemented using feedback control techniques to regulate their aggregate power. To achieve high performance, such feedback control techniques require an accurate mathematical model of the TCL aggregate dynamics. Although such models have been developed, they appear too complex to be effectively used in control design. In this paper we develop a mathematical model aimed at the design of a model-based feedback control strategy. The proposed model analytically characterises the aggregate power response of a population of ACs to a simultaneous step change in temperature set points. Based on this model, we then derive, and completely parametrise in terms of the ACs ensemble properties, a reduced-order mathematical model to design an internal-model controller that regulates aggregate power by broadcasting temperature set-point offset changes. The proposed controller achieves high LM performance provided the ACs are equipped with high resolution thermostats. With coarser resolution thermostats, which are typical in present commercial and residential ACs, performance deteriorates significantly. This limitation is overcome by subdividing the population into clusters of ACs that receive a coarse-grained, cluster-dependent control signal. The proposed clustering technique recovers the performance achieved with high resolution thermostats at the expense of some additional comfort penalty, which can be quantified using the controller output.

Crown Copyright © 2011 Published by Elsevier Ltd. All rights reserved.

1. Introduction

Demand side management (DSM) technologies include a wide range of strategies to reduce power usage to balance supply and demand in electricity markets at peak periods, reducing the pressure for upgrades in power generation and distribution infrastructure. Studies by the International Energy Agency (IEA) suggest that DSM is more cost-effective and sustainable than conventional policies based on supply side: each \$1 invested in DSM has been estimated to offset \$2 spent in supply side improvements, while contributing to reduce greenhouse emissions [1, Chapters 7, 8]. Load management (LM) is a DSM strategy that aims to balance supply and demand by reducing the power use of electrical devices during critical periods instead of increasing the power generation. As smart meters and appliances slowly become mainstream, LM

technologies gain strength as an alternative for the electricity market [2,3].

An important class of electric loads that can be integrated in LM strategies with full responsiveness are thermostatically controlled loads (TCLs) [3,4]. TCLs encompass devices such as air conditioners (ACs), fridges, and space and water heaters, which are typically responsible for a large proportion of the residential energy demand [5]. The flexibility of TCLs for demand control comes as a result of their thermal inertia: TCLs may be viewed as a distributed energy storage resource that can be controlled with constraints imposed by acceptable impact on end users.

The operation of TCLs can be manipulated by LM for various reasons, the most common one being to reduce the power demand during periods of high electricity prices or high electricity demand. More recently, electricity markets are starting to consider the participation of loads alongside conventional supply-side resources [3,6], a trend that is likely to become more prominent as aggregations of TCLs become more competitive with traditional supply side approaches [7–9]. Another application of TCLs for LM is to control

* Corresponding author at: CSIRO Energy Technology, 10 Murray Dwyer Circuit, Mayfield West, NSW 2304, Australia.

E-mail address: cristian.perfumo@csiro.au (C. Perfumo).

demand to track variations in generated energy, such as that typically observed in renewable generation [4].

The aggregate power requirement of a population of TCLs may be effectively controlled by manipulating a common temperature set point offset. As shown in [4,5], small common offset changes in temperature set point may be broadcast to the population to control their aggregate power with minimal impact to individual users.

However, common set point changes in a large population of TCLs may produce large undesirable transients in their collective power response, as their states will tend to synchronise under a common disturbance. Such transients are typically observed in TCLs when the power supply is interrupted for an extended period of time and then simultaneously restored, a phenomenon traditionally known as cold-load pickup [10]. Similar undesirable collective behaviour may even occur in large populations of loads under randomised (but uncoordinated) autonomous control, as illustrated in [3] for populations of plug-in electric vehicles.

In the present paper we consider the design of a model-based, coordinated feedback strategy to control the aggregate power demand of ACs by manipulating a common offset in their temperature set points. Using modern control techniques, system dynamics may be controlled to high degrees of performance by incorporating an accurate model of the uncontrolled dynamics in the design of the controller, a notion referred to as *internal model control* [11]. A core contribution of this paper is the development of a model that accurately captures the collective dynamics of TCLs while maintaining a mathematical complexity suitable for an effective design of the model-based controller.

Substantial research on modelling populations of TCLs, and using these models, is available. Some authors concentrate on models based on first principles [12,10,13–17], while others focus on identifying the model parameters from a real population of devices [18,19], including the use of black-box model identification techniques [4,20]. However, the majority of these models are too complex to be effectively used by well-understood feedback control design methods.

For example, a well-known model developed by Malhame and Chong (M&C) [14,15] consists of a set of Fokker-Planck diffusion equations describing the probability density distribution of temperature in a population of identical heaters that condition identical spaces. By integrating the temperature distribution, the probability of a device being operating at a certain time (and therefore the power usage profile of the population at that time) can be calculated. Using M&C's model in the design of a feedback control algorithm is very difficult because this model is formulated as a system of partial differential equations with no closed-form solution in the general case. This difficulty appears as a main cause for most of the existing approaches to control a population of TLCs for LM to be in open loop [21,17,5].

One exception to open-loop control of TCLs is the approach proposed by Callaway in [4]. Callaway proposes the use of a broadcast temperature set point offset as the output signal of a feedback-controller to regulate the aggregate power output of a population of ACs. Core dynamics of the model in [14] are captured in [4] by a first-order linear ARMAX (AutoRegressive Moving Average eXogenous) model obtained using standard black-box system identification techniques [22]. This identified model is then used in [4] to develop a *minimum variance controller* that drives the power output of the population to track the generated power of a wind farm.

The idea in [4] of using a feedback controller for the aggregate power demand of a population of TCLs using a global temperature set point offset as the control signal has recently been adopted in [23,24]. In [23], the authors incorporate a control signal to M&C's model and develop a Lyapunov-stable algorithm to control a population of homogeneous devices and rooms. In [24] a homogeneous population is also considered to obtain an initial undamped model

to which the authors manually add a damping coefficient to adjust the modelled response to the simulated one. The damped model is then used to develop a linear quadratic regulator controller.

Other types of control signal could be used to control the aggregate power output (such as toggling the ON/OFF state of all the devices in certain temperature range [25]). However, the benefit of a temperature set point global offset is that it directly relates to user comfort. This relation easily allows the controlling entity (presumably the electricity utility) to estimate the end user impact in a LM scenario. One limitation of the control approach in [4,23,24] is that it relies on temperature set point changes as little as 0.0025 °C, whereas the typical set point resolution of hardware installed nowadays is two orders of magnitude higher, at around 0.1 to 0.5 °C.

The present paper formulates an aggregated model based on first principles that represents, under certain assumptions, the power consumption of a population of ACs over time. We use this model to show that the power response of a group of ACs to a temperature set point common change presents underdamped oscillations, a phenomenon widely known and simulated but, to the best of our knowledge, not before characterised analytically. From our analysis, we derive a linear time-invariant (LTI) second order mathematical model that we use to design an internal-model controller to regulate aggregate power response by broadcasting set point temperature offsets. All the parameters of the derived second order model are analytically characterised by the ensemble properties of the population of ACs considered. Thus, the main characteristics of the aggregate power response, and a model-based control strategy, are directly parameterised by the properties of the population of ACs. This is the main original contribution of the paper to modelling and control of TCLs.

The simulation results show accurate load control performance, provided the temperature resolution of the ACs is fine enough (similarly to [4]). As an implementation of the proposed control approach for coarse-resolution ACs (such as the ones installed currently at most homes), we subdivide the population into clusters that receive cluster-dependant coarse-resolution control signals following a method similar to that of *demultiplexing* in communication networks. Thus, spatial diversification is exploited to reduce the quantisation constraints on the aggregate control signal. We show that the proposed clustering technique recovers the control performance of a common finely quantised control signal at the expense of minor comfort penalties. Such penalties can be quantified by estimating the predicted percentage of dissatisfied people (PPD) [26] based on the temperature set point offsets used as the control signal.

The remainder of the paper is organised as follows: Section 2 presents a well-known single-AC model and describes how we use it to simulate populations of devices under a LM scenario. Section 3 develops the proposed model of a population of ACs in three passes, in each of which we decrease the complexity of the model by incorporating additional assumptions and simplifying approximations. We end up in a fully parameterised second-order LTI model. Section 4 uses this LTI model to design the proposed model-based controller for LM of the population, and Section 5 refines this controller for implementation on coarse-resolution thermostats. Section 6 shows how to quantify the comfort impact in the proposed LM scenario using the control signals broadcast to the clusters. Finally, Section 7 summarises our conclusions.

2. A preliminary model for the power dynamics of an aggregation of ACs

We consider a population of ACs, each of which regulates the average temperature $\theta(t)$ of a room by means of a thermostat and relay actuator with state $m(t) \in \{0,1\}$, which determines whether the compressor is switched on ($m = 1$) or off ($m = 0$)

according to a pre-specified hysteresis band $[\theta_-, \theta_+]$ and a desired temperature set point $(\theta_- + \theta_+)/2$. The dynamic behaviour of such an AC can be more precisely described by the two-state hybrid mathematical model

$$\frac{d\theta(t)}{dt} = -\frac{1}{CR}[\theta(t) - \theta_a + m(t)RP + w(t)], \quad (1)$$

$$m(t^+) = \begin{cases} 0 & \text{if } \theta(t) \leq \theta_- + u(t) \\ 1 & \text{if } \theta(t) \geq \theta_+ + u(t) \\ m(t) & \text{otherwise,} \end{cases} \quad (2)$$

where θ_a represents the ambient temperature (assumed constant), C and R are the thermal capacitance (kWh/°C) and thermal resistance (°C/kW) of the room being air-conditioned, and P (kW) is the thermal power of the AC, which is the electrical power times its coefficient of performance. The term $w(t)$ in (1) represents unpredictable thermal disturbances (heat gains or losses) such as the effect of people in the rooms, open doors and windows, and appliances. In (2), $u(t)$ represents a temperature set point offset that can be manipulated over time.

The hybrid model (1), (2) has been used (with $u(t) = 0$) in early mathematical studies of the dynamics of aggregate power of populations of ACs [12,10], and more recently in [4], which incorporated the temperature set point offset $u(t)$ as an external control signal to achieve load tracking to compensate the variability of wind energy generators. Note that controlling a group of ACs using a set point offset avoids forcing a global absolute reference temperature for all the devices, allowing the occupants to choose the temperature set point they are most comfortable with.

To describe the collective demand of a population of n ACs, let $i \in [1, 2, \dots, n]$ denote the index representing the i th device, which is described by a hybrid model of the form (1) and (2). Then, the evolution of the aggregate normalised power demand $D(t)$ (normalised by the maximum power demand of the population) is given by the ratio

$$D(t) = \frac{\sum_{i=1}^n m_i(t)D_i}{\sum_{i=1}^n D_i}, \quad (3)$$

where $m_i(t)$ represents the discrete state and D_i the electrical power (P_i divided by the coefficient of performance) of the i th AC in the population.

The aggregation in (3) suggests a simple way to numerically simulate the evolution of the aggregate power of such population of ACs by sampling values for the constant parameters C , R and P for each AC from some pre-specified distributions and then computing and adding together the individual power outputs [4]. Thus, the single-device model described by (1) and (2) has been extensively used as a starting point for numerical studies of the behaviour of populations of TCLs [12,10,14,16,17,4].

In contrast, a mathematically precise analysis of (3) is challenging, since it comprises a distribution of n (potentially thousands) non-linear dynamic systems, each of which evolves independently according to (1) and (2). One of the earliest and most important contributions to such analysis is the work by Malhame and Chong [14], who developed a system of Fokker-Planck equations to describe the time evolution of the probability distribution of an AC being at certain temperature. The recent work by Callaway [4] refined the model developed in [14] and further simplified it for the purposes of control design for load tracking of variable renewable generation.

The present paper uses the model in (1) and (2) both as a basis of developing a simplified analytical model describing the aggregate power response of an entire population of ACs, and to perform numerical simulations of large numbers of ACs. The latter are used to validate the proposed simplified analytical model, presented in the following section.

3. Theoretical analysis towards a simple population model based on first principles

In this section we present an expression that describes how the proportion of ACs that are switched on changes over time when a population of devices is excited with a step set point offset signal. More importantly, we use this expression to derive a simplified model, suitable for use in control design.

We start by making the following simplifying assumptions about the continuous state of the single-AC model in (1):

- H.1 All of the ACs in the population have the same set point temperature $\theta_{ref} = (\theta_+ + \theta_-)/2$ and the same hysteresis width $\theta_+ - \theta_- = 1$.
- H.2 At $t = 0$, the temperatures are uniformly distributed in the interval $[\theta_-, \theta_+]$.
- H.3 The parameter C is distributed in the population according to some probability distribution. The parameter R is the same for all the ACs in the population, and the same applies to P .
- H.4 For any AC, the temperature decreases when the AC is operating at the same rate it rises when the device is switched off, which implies that $RP = 2(\theta_a - \theta_{ref})$ for each device.
- H.5 $\theta_+ - \theta_- \ll |\theta_a - \theta_{ref}|$, i.e., for each AC, the rate at which the temperature changes is constant, so that the temperature describes the triangular waveform as shown at the top of Fig. 2.
- H.6 The noise term $w(t)$ for each AC is negligible.

Assumptions H.1, H.3, H.4 and H.6 are required simplifications for our analysis. While they may seem overly restrictive, in Section 5 we show that the controller designed using these simplifying assumptions still can preserve good performance even when these assumptions are relaxed.

H.4 implies that the duty cycle in steady state is 50%. H.2 is reasonable since the temperatures of a group of devices that have been running independently for long enough are distributed almost uniformly [4]. It is also sensible to make assumption H.5 as the hysteresis width is expected to be significantly smaller than the difference between the ambient and reference temperatures, rendering rates of temperature change constant rather than variable (as obtained when (1) is solved [23]).

Finally, we assume H.6 because, when analysing a population of ACs as a whole, the variability that $w(t)$ introduces is small compared to the one introduced by the fact that the parameter C is, in general, different for each AC.

We analyse how a population of ACs, each governed by (1), (2) and H.1–H.6, reacts when a step change of amplitude 0.5 is applied at $t = 0$, shifting the hysteresis boundaries to the right. We will refer to these new boundaries as $\theta_-^{post} = \theta_- + 0.5$ and $\theta_+^{post} = \theta_+ + 0.5$. Fig. 1 shows the temperature distributions just before and after the step. The top part of each subfigure represents the ACs that are operating whereas the bottom part depicts the ones that are switched off. The arrows indicate in which direction the temperatures are moving. We can see in Fig. 1a and b that the ACs that were operating and had temperatures in $[\theta_-, \theta_-^{post}]$ before the step switched off after it.

Under assumptions H.1–H.6, the absolute value of the rate at which the temperature $\theta_i(t)$ changes for the i th AC (as described in (1)) is given by the constant v_i , defined by

$$\left| \frac{d\theta_i(t)}{dt} \right| \approx v_i = \frac{\theta_a - \theta_{ref}}{C_i R}. \quad (4)$$

We introduce the variable

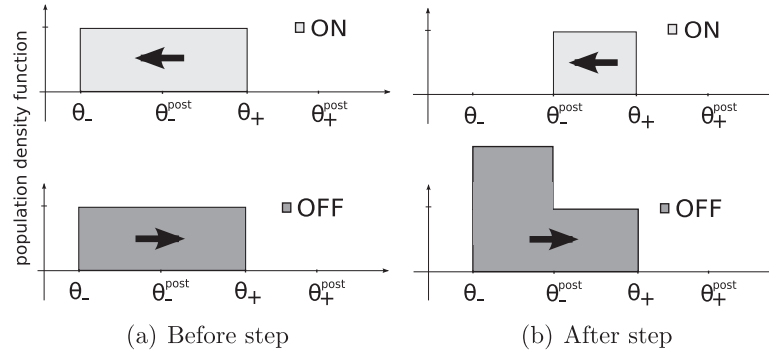


Fig. 1. Distribution of temperatures before and after a 0.5 °C step in the set point.

$$x_i(t) = x_i^0 + v_i t, \quad (5)$$

where

$$x_i^0 = \begin{cases} 1 + \theta_i(0) - \theta_-^{post} & \text{if } \frac{d\theta_i}{dt}(0^-) > 0, \\ \theta_+^{post} - \theta_i(0) & \text{if } \frac{d\theta_i}{dt}(0^-) < 0. \end{cases} \quad (6)$$

Intuitively, we can say that $x_i(t)$ is the “unwrapped” mapping of the temperature $\theta_i(t)$. Fig. 2 illustrates the equivalence between $\theta_i(t)$ and $x_i(t)$, the latter being a straight line, as implied by H.4 and H.5. Note that in Fig. 2, at the times the temperature $\theta_i(t)$ changes direction, $x_i(t)$ reaches an integer value, changing from an even (on) to an odd (off) interval, or vice versa. We refer to an interval of values of $x_i(t)$ as *odd* whenever $x_i(t) \in [2k+1, 2k+2]$, for $k = 0, 1, \dots$, and *even* whenever $x_i(t) \in [2k, 2k+1]$.

Now we are ready to formalise the step response of a population of ACs in the following proposition (its proof can be found in Appendix A.1).

Proposition 1. For a population of ACs where the dynamics of each device are described by (1) and (2), under the assumptions H.1–H.6, if the temperature set point of all the devices in the population is raised by 0.5 °C at time $t = 0$, the probability $D(t)$ that a randomly-selected AC be operating at time t is given by:

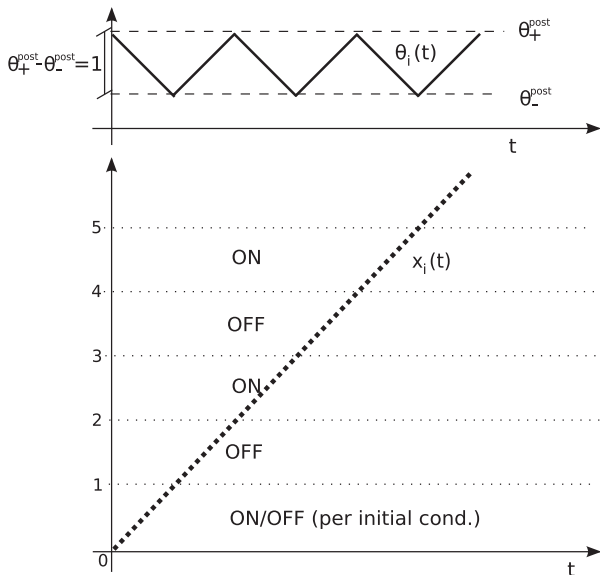


Fig. 2. Top: temperature $\theta_i(t)$ of one AC. Bottom: its “unwrapped” equivalent $x_i(t)$.

$$D(t) = \frac{\Pr[x(t) < 1]}{3} + \sum_{k=1}^{\infty} \Pr[x(t) < 2k+1] - \Pr[x(t) < 2k] \quad (7)$$

where $\Pr[\cdot]$ is the probability operator.

Proposition 1 says that the probability that one randomly-picked AC be operating at time t is given by the probability that $x(t)$ is in an even interval plus a correction for the initial condition arising from H.2. Note that for a large enough population, the probability (7) is equivalent to the proportion of ACs in the population operating at time t . Also, because all of the ACs have the same power (H.3), this proportion is equal to the power consumption of all of the devices operating at time t , normalised by the maximum demand (all of the ACs in the population operating), as described in (3). Therefore, throughout this paper, we deliberately use the notation $D(t)$ to refer to the probability of a randomly-selected AC operating, the proportion of operating ACs in the population and the normalised power consumption.

In order to calculate actual values for (7), it is necessary to know how the temperatures change for each AC in the population. This can be described by characterising the distribution of each of the parameters R , P and C in the population. As in [4], we adopt log-normal distributions for these parameters as they are suitable for non-negative parameters and have a complexity of description that is only moderate.

Assuming that all of the ACs have the same thermal resistance R and the same thermal power P , and the thermal capacitance C is distributed log-normally in the population, Eq. (7) can be more explicitly calculated, as shown in the following corollary of Proposition 1.

Corollary 1. Under the assumptions of Proposition 1, let all of the ACs have the same thermal resistance R and the same thermal power P , and let the thermal capacitance C be distributed log-normally with mean μ_C and standard deviation σ_C . Then the speed v at which the temperature changes (Eq. (4)) is distributed log-normally and the ratio σ_{rel} between its standard deviation σ_v and mean μ_v is

$$\sigma_{rel} = \frac{\sigma_v}{\mu_v} = \frac{\sigma_C}{\mu_C}. \quad (8)$$

Furthermore, the probability $D(t)$ that a randomly picked AC be operating at time t can then be approximated by

$$D(t) \approx \frac{1}{6} + \frac{1}{6} \operatorname{erf} \left[\frac{\log(1) - \log(\mu_x(0) + \mu_v t)}{\sqrt{2} \sigma_{rel}} \right] + \frac{1}{2} \times \sum_{k=2}^{\infty} (-1)^{k+1} \operatorname{erf} \left[\frac{\log(k) - \log(\mu_x(0) + \mu_v t)}{\sqrt{2} \sigma_{rel}} \right]. \quad (9)$$

where $\operatorname{erf}[\cdot]$ is the Gauss error function and $\mu_x(t)$ is the mean of the values x at time t .

A formal proof of Corollary 1 is presented in Appendix A.2.

Fig. 3 plots in dashed lines the expected power response according to (9), for different values of σ_{rel} , to a 0.5°C step at $t = 0$. The output is normalised to the maximum power output (all ACs turned on). The figure also shows, as a solid black line, the output to the same input when we simulate using the PowerDEVs tool [27], 10,000 ACs (according to (3)) assuming homogeneous thermal resistance R thermal power P . The thermal capacitance C is distributed log-normally in the population with a standard deviation to mean ratio of σ_{rel} . The values for the parameters used for the simulations, detailed in Table 1, are the same as in [4], except for the ambient temperature, which was adjusted to obtain a duty cycle of 0.5, to make a comparison with the theoretical results possible. Fig. 3 validates Proposition 1 and Corollary 1, as the analytical and simulated responses are very close.

Note that in the simulated results, H.4–H.6 were relaxed, considering a much more realistic scenario. In particular, it is important to note that relaxing H.4 implies that, in general terms, each room gets cooled down at a potentially different rate than it heats up (no fixed relation between R and P). The aspect of H.4 that still holds for the simulated results is that the average duty cycle in the population is 0.5, which is why the simulated results present a steady state normalised power consumption of 0.5. The authors are currently working on an extension of the analysis presented in this section to contemplate different duty cycles, which is a more realistic scenario as, for example, a duty cycle higher than 0.5 would be expected for very hot days, and a lower one for cooler days (especially for populations of oversized ACs, which are common). Nevertheless, in Section 5 we show that the controller designed using the model obtained in the present section, assuming a duty cycle of 0.5, shows robustness even when used on a plant with an average duty cycle far from 0.5.

When the values of all three parameters R , P and C are distributed over a range within the population, there is more variability in the speed at which the temperature changes, and therefore the ACs tend to desynchronise faster (the aggregate responses simulated with the

Table 1
Simulation parameters.

Parameter	Value	Description
R	$2^\circ\text{C}/\text{kW}$	Mean thermal resistance
C	$10 \text{ kWh}/^\circ\text{C}$	Mean thermal capacitance
P	14 kW	Mean thermal power
θ_-	19.5°C	Lower end of hysteresis band
θ_+	20.5°C	Higher end of hysteresis band
θ_a	32°C	Ambient temperature
σ_w	0.01	Standard deviation of the noise process w in Eq. (1)
σ_{rel}	0.05/0.1/0.2/0.5	Standard deviation of log-normal distributions as a fraction of the mean value for R , C and P

parameters from Table 1, not shown, are more damped than those in Fig. 3). Nevertheless, the dominant dynamics are the same for both cases: the power response of a population of ACs to a step in the temperature set point presents damped oscillations.

The expression (9) can be used to analytically show that the transients in aggregate power due to a step change in temperature set point are characterised by underdamped oscillations. The following proposition formalises this observation.

Proposition 2. Under the assumptions of Corollary 1, the initial transients in the response of the aggregate power to a step change in temperature set point displays oscillations with period

$$T \approx 2/\mu_v, \quad (10)$$

and peaks with amplitudes $A(t)$ that decay with time according to the approximate envelope bound

$$1 - \text{erf}(1/z(t)) \leq A(t) \leq \text{erf}(1/z(t)), \quad (11)$$

where

$$z(t) = 2\sqrt{2}\sigma_{rel}(\mu_x(0) + \mu_v t - 1/2). \quad (12)$$

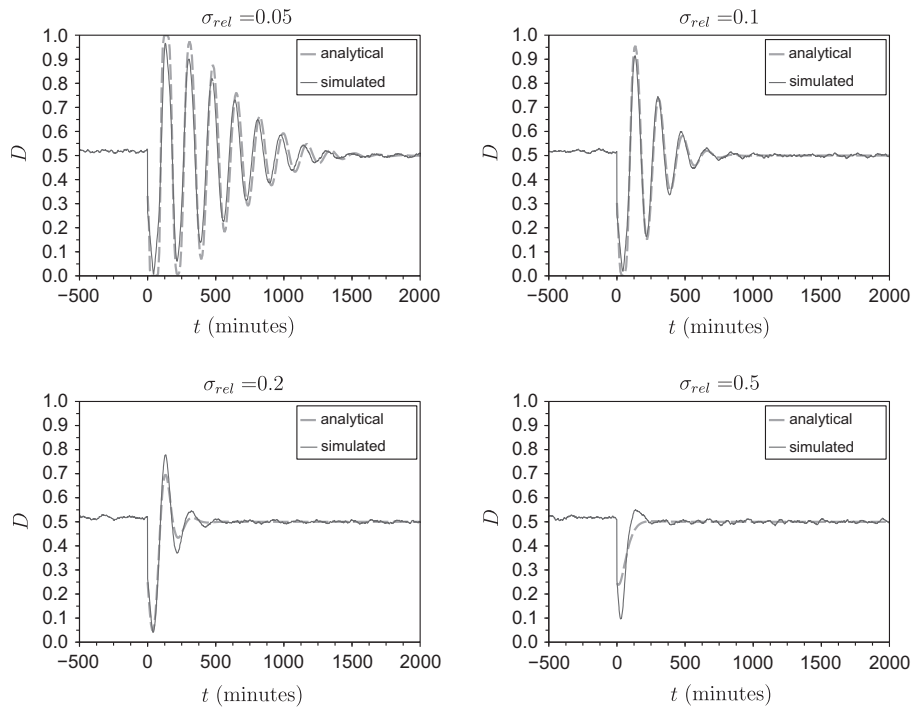


Fig. 3. For different values of σ_{rel} , expected and simulated D (normalised power demand) response to an increase of 0.5°C in all ACs' temperature set points. The simulated results are for 10,000 ACs with log-normally distributed C and constant R and P , and were obtained using the parameters from Table 1. To obtain a duty cycle of 0.5, θ_a was set to $\theta_{ref} + RP/2 = 34^\circ\text{C}$.

A formal proof of Proposition 2 can be found in Appendix A.3.

Fig. 4 depicts the envelope bound (11) of the oscillations in the step power response along with the expected power output calculated from (9) for different values of σ_{rel} as a function of the offset and rescaled time z , as defined in (12), which allows us to plot all responses within a common envelope. Each curve starts at the corresponding value of z for $t = 0$. We can see in Fig. 4 that for a range of values of σ_{rel} , the bounds in (11) capture accurately the envelope of the response. We also observe in Fig. 4 that for values of z larger than 1.8, the bounds in Eq. (11) are no longer valid (see proof of Proposition 2). Nevertheless, the fundamental dynamics of the response (i.e., damped oscillations) are still observed for larger values of z , and hence t .

The analysis carried out so far demonstrates that the step response of the aggregate power of a population of ACs under the assumptions H.1–H.6 is dominated by decaying oscillations, which corroborates the simulation results reported by a number of authors [10,15,13,4]. In particular, if the duty cycle is around 0.5, the “mountains” and “valleys” of $D(t)$ have equal semi-periods, which suggests that second order dynamics would be a suitable approximation of the aggregate power response. The following corollary to Proposition 2 provides formulas to compute the characteristic polynomial of such second order dynamics directly from the parameters describing the population of ACs (i.e. Table 1).

Corollary 2. Under the assumptions of Proposition 2, the response of a population of ACs to a 0.5 °C step can be modelled as

$$D(t) = D_{ss}(\theta_{ref}) + \mathcal{L}^{-1}\{G_p(s)0.5/s\} \quad (13)$$

where $G_p(s)$ is a second order linear, time-invariant (LTI) system characterised by the transfer function

$$G_p(s) = \frac{b_2 s^2 + b_1 s + b_0}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (14)$$

whose parameters are

$$\begin{aligned} \zeta &= \frac{\log(r)}{\sqrt{\pi^2 + \log^2(r)}}; & \omega_n &= \frac{\pi\mu_v}{\sqrt{1 - \zeta^2}}; \\ b_0 &= \frac{\omega_n^2(D_{ss}(\theta_{ref}) + 0.5) - D_{ss}(\theta_{ref})}{0.5}; \\ b_1 &= 0.5\mu_v + 2b_2\zeta\omega_n \quad \text{and} \quad b_2 = D_{ss}(\theta_{ref}). \end{aligned} \quad (15)$$

where




Fig. 4. Envelope of the power peaks. The dotted and dashed lines represent the normalised power demand D for different values of σ_{rel} according to (9), all of them adjusted to the offset and rescaled time variable z . The continuous line is the envelope given by (11) and (12).

$$r = \frac{|\operatorname{erf}\left(\frac{1}{0.9+2\sqrt{2}\sigma_{rel}}\right) - 0.5|}{|\operatorname{erf}\left(\frac{1}{0.9}\right) - 0.5|}$$

and

$$D_{ss}(\theta) = \left(1 + \frac{\log\left(1 + \frac{H}{\theta_a - \theta - H/2}\right)}{\log\left(1 + \frac{H}{PR + \theta - \theta_a - H/2}\right)}\right)^{-1} \quad (16)$$

and

$$H = \theta_+ - \theta_1$$

Corollary 2 (whose proof can be found in Appendix A.4) provides direct, rule-of-thumb, information about the step response of a population of ACs. For example, using Corollary 2 we compute $\zeta = 0.259$ and $\omega_n = 0.033$ for the scenario shown in the bottom-left plot in Fig. 3. These values of ζ and ω_n give an estimated 1% settling time $t_s = 4.6/(\zeta\omega_n) = 525$ (min) [28], which means that, for the considered scenario, the step response will approximately decay to 1% in 525 minutes. By comparison with the simulated response of 10,000 ACs shown at the bottom-left plot in Fig. 3, we can see that the settling time computed using Corollary 2 is an excellent estimate of the actual settling time of the collective response. Note that to compute such an estimate otherwise, we would need to perform one of these two computationally intensive tasks: (a) run numerical simulations of a large number of ACs (such as the ones shown in solid black in Fig. 3) or (b) approximate numerically M&C's set of Fokker-Planck differential equations [14].

More importantly, knowing the transfer function (14) allows us to design a model-based controller for the a population of ACs based only on the physical parameters of such a population. Such a controller is presented in the following section.

4. Controlling the plant

In this section we present the design of a controller for a population of ACs that is able to adjust the power output to a desired profile by manipulating the temperature set points of the devices. We use a model-based control structure known as internal model control (IMC) [11]. Such a model-based control structure incorporates a model of the plant as a fundamental part of the controller. The following three subsections describe the model of the plant for control design purposes, its implementation in the design of the controller, and present the closed-loop responses obtained under different simulated scenarios.

4.1. A linear, time-invariant model for a population of ACs for control design

We use the second order linear, time-invariant (LTI) model structure in (14) as a model for the aggregate response.

Table 2 shows the values of (15) obtained for a population of ACs described by the parameters in Table 1 ($\sigma_{rel} = 0.2$). This combination of parameters of the population is the same as in [4], except for the hysteresis width, for which we use 1 °C as opposed to 0.5 °C, as we consider 0.5 °C overly tight.

We also show in Table 2 the parameters obtained when a second order model is manually identified from the simulated response to a 0.5 °C step of a population of 10,000 ACs as described

Table 2

Transfer function parameters of the second order model in (14) calculated from (15) and identified manually.

Method	ζ	ω_n	b_0	b_1	b_2
As per (15)	0.259	0.033	3.70×10^{-5}	0.029	0.446
Manually identified	0.3505	0.0326	3.61×10^{-5}	0.038	0.45

in (3) with parameters as per Table 1 ($\sigma_{rel} = 0.2$). We fit the parameters of the proposed second order model structure by measuring peaks, period and steady state gain in the simulated step response of 10,000 ACs.

Fig. 5 compares the simulated response with that of the calculated second order LTI model using (15), and the one manually identified. The figure shows that the calculated model can capture the dominant dynamics of the simulated response very closely for the considered set of parameters. Moreover, from Table 2 and Fig. 5 we can conclude that the second order model calculated from the parameters of the population of ACs is comparable to a manually identified model. Thus, throughout the reminder of this paper we will use the calculated model using (15) for our control design.

Note that, as per Table 1, we simulate a population where the parameters P, R and C are *all* distributed lognormally ($\sigma_{rel} = 0.2$




Fig. 5. Step power response D for 10,000 simulated ACs with parameters described in Table 1 ($\sigma_{rel} = 0.2$) and using the second order systems calculated with (15) and identified manually.




Fig. 6. Internal model control structure.

indicates that the standard deviation of each parameter is calculated as its mean times 0.2), which relaxes assumptions H.3 and H.4. In fact, all the simulated results that we present in the reminder of this paper are obtained from populations where P, R and C are all distributed. We do this to illustrate the robustness of the proposed control algorithm, which preserves the desired performance even when we move away from the simplifying assumptions on which the controller is based.

4.2. Using the calculated LTI model to control the population via internal model control

Fig. 6 shows a typical IMC structure. The input signal $D_{ref}(t)$ to the controller represents the desired normalised aggregate power (proportion) of operating ACs in the population at time t . Given this reference signal and the actual normalised aggregate power $D(t)$ observed in the population, the controller computes the temperature set point offset signal $u(t)$ to be broadcast to the ACs.

The IMC controller encompasses a model of the plant $G_p(s)$, its inverse (or a stable approximation if the inverse does not exist) $G_p^{-1}(s)$, and the target desired behaviour in closed loop represented by its transfer function $G_d(s)$. For simplicity, we propose

$$G_d(s) = \frac{1}{T_c s + 1}$$

where T_c is the desired closed-loop time constant of the system. For the results presented in this paper we chose $T_c = 2$ (min), since it allows reaching the steady state reference value within a reasonable time (10 min) using an affordable amount of control effort.

4.3. Preliminary closed-loop performance and implementation issues

We propose a LM scenario with two stages: power reduction and comfort recovery. During power reduction, the controller manipulates the temperature set points to achieve an aggregate power output lower than the starting steady state value. In recovery, the demand is raised over the steady state value to make up for the increase in the temperatures of the rooms caused during the reduction period.

An instance of this scenario is presented in Fig. 7a. The ACs run uncoordinated (global control $u(t) = 0$) for the first 100 minutes, with a steady state power output of 44%, relative to the maximum. The reduction period takes place between $t = 100$ and $t = 340$ (min), reducing the power output by 10% of the maximum load.

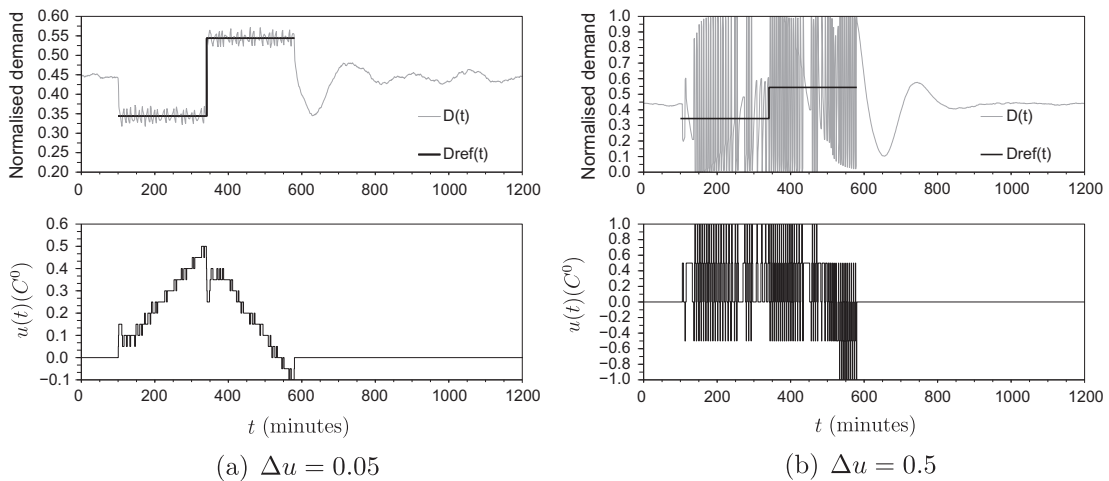


Fig. 7. LM scenario on a population of 10,000 ACs between $t = 100$ and $t = 580$ (min). The control aims to make the normalised power output $D(t)$ follow the reference signal $D_{ref}(t)$ (top plot). The control signal $u(t)$ is shown in the bottom plots. The quantisation of $u(t)$ is $\Delta u = 0.05$ for (a) and $\Delta u = 0.5$ for (b).

The recovery takes place between $t = 340$ and $t = 580$ (min), when the power output is maintained at 54% (the steady state output plus 10% of the maximum). The LM scenario finalises at $t = 580$, when the ACs continue to operate without global control.

This scenario is far from optimal for LM. The oscillations in D observed after $t = 580$ indicate that, once the devices stop being controlled, the system is not left in steady state. One could control the system after $t = 580$ in order to gradually bring D back to steady state, but it should be kept in mind that the control period should be finite, as it is unlikely for a customer to accept that their ACs be manipulated indefinitely. It is not the objective of the present paper to explore optimal recovery phases but rather to show a way to control a population of ACs to follow a desired reference signal.

Fig. 7b presents the same scenario as Fig. 7a except for a coarser granularity of the control signal. Because the temperature sensors in the ACs do not have infinite resolution and accuracy, the input signal must be quantised for the control to be implementable. We observe in Fig. 7b that for a coarse quantisation $\Delta u = 0.5$, the controller is unable to track the reference since, as we saw in Fig. 3, a step of 0.5°C implies a large (transient) change in the output. Conversely, when we use a granularity small enough $\Delta u = 0.05$, the controller successfully makes the plant follow the reference signal (Fig. 7a). Thus, the granularity of the signal plays a key role in the quality of the output.

From the results presented in Fig. 7 arises the following implementation difficulty: how can LM be implemented using a temperature set point change as the control signal in a feedback-controlled system, if small values of Δu are needed for an acceptable response (as also suggested in [4]) and, on the other hand, the set point resolution of the majority of residential ACs is 0.1 , 0.5 or even 1°C ? In the following section we propose a practical solution to this issue.

5. Cluster-based control implementation: exploiting spatial diversification to solve the problem of controlling currently installed ACs

To resolve the temperature set point resolution issue pointed out in the previous section, we propose to divide the population of ACs into clusters, and implement a mapper that demultiplexes the fine-granularity control signal for the entire population into multiple coarse-granularity control signals, one for each cluster defined. Fig. 8 presents a schematic diagram of how the mapper interacts with the clusters and the controller. Note that no changes in the controller are needed, since the mapper is considered part of the plant. The mapper receives the “global” control signal $u(t)$, which is discretised very finely to a small quantum Δu (e.g. 0.05°C) and generates L different “cluster” signals $u_i(t)$, which are multiples of a coarse quantisation $\Delta' u$ that the ACs can deal with (e.g. 0.5°C). If we choose $L\Delta u$ to be a multiple of $\Delta' u$, each cluster signal is defined¹ by:

$$u_i(t) = \begin{cases} \text{floor}\left[\frac{u(t)}{\Delta' u}\right] \Delta' u & \text{if } i > L \frac{\text{mod}(u(t), \Delta' u)}{\Delta' u} \\ \text{floor}\left[\frac{u(t)}{\Delta' u}\right] \Delta' u + \Delta' u & \text{if } i \leq L \frac{\text{mod}(u(t), \Delta' u)}{\Delta' u} \end{cases} \quad (17)$$

For example, if the population is divided in 10 clusters and $u(t_m) = 1.15$ at $t = t_m$, then the cluster control signals will be $u_i(t_m) = 1.5$ for $i = 1, 2, 3$ and $u_j(t_m) = 1.0$ for $j = 4, \dots, 10$.

Fig. 9 shows the same LM scenario as Fig. 7: 10,000 ACs being controlled to reduce the load from 44% to 34% between $t = 100$




Fig. 8. Redefinition of the plant when incorporating a mapper to translate the global control signal to several different signals, one per cluster.




Fig. 9. Normalised aggregate power demand D (top) and global and cluster-level input control signals $u(t)$, $u_1(t)$, $u_{10}(t)$ (bottom) for a LM scenario between $t = 100$ and $t = 580$ (min) on a population of 10,000 ACs. The control aims to make the normalised power output $D(t)$ follow the reference signal $D_{ref}(t)$. The 10,000 ACs are divided into 10 clusters of 1000 ACs each. The bottom plot shows the global control signal (solid) as well as the individual signals for clusters 1 and 10 (dotted and dashed respectively).

and $t = 340$ (min) and to increase it to 54% between $t = 340$ and $t = 580$. In Fig. 9 the population is divided in 10 clusters of 1000 ACs each, using (17) to generate the 10 signals. We use a fine-grain $\Delta u = 0.05^\circ\text{C}$ and coarse-grain $\Delta' u = 0.5^\circ\text{C}$. The bottom plot in Fig. 9 shows the global control signal $u(t)$ as well as the individual signals for the first and tenth clusters $u_1(t)$ and $u_{10}(t)$. From (17), it follows that the rest of the cluster control signals ($u_2(t)$, $u_3(t)$, \dots , $u_9(t)$) are bounded above by $u_1(t)$ and below by $u_{10}(t)$. Also, note that at any time t , $|u(t) - u_i(t)| < \Delta' u$ and $|u_i(t) - u_j(t)| \in \{0, \Delta' u\}$, which is to say that the distance between any cluster control signal and the global control signal is less than the quantisation $\Delta' u$, and the distance between two individual signals is either 0 or $\Delta' u$.

Clustering does not hurt output performance, as the variance of the difference between the reference and the output in Fig. 9 is slightly smaller than that of Fig. 7a (no clusters), for both the reduction and recovery phases. However, there is a comfort penalty associated with clustering, as in the example in Fig. 9 each AC receives a set point change up to ten times larger than that of the scenario in Fig. 7a. We quantify this impact on comfort in Section 6.

In Fig. 7a and Fig. 9, we can see that our controller performs well even when the parameters P , R and C are distributed in the population, relaxing assumption H.4. Fig. 10 shows a scenario that includes this relaxation and also considers a duty cycle far from the 50% implied by H.4. We simulate such a scenario with the same parameters as that from Fig. 9, except for the ambient temperature, which is set to 45°C , causing a duty cycle in steady state of 85%.

¹ Fairness problems such as the one in (17), where the clusters with lower numbers always receive a larger temperature set point offset, are easily overcome by assigning virtual numbers to the clusters and changing these numbers every time a LM scenario takes place.




Fig. 10. Normalised output power demand D (top) and global input control signal $u_i(t)$ (bottom) for a LM scenario between $t = 100$ and $t = 580$ (min) on a population of 10,000 ACs with 85% mean duty cycle. The control aims to make the normalised power output $D(t)$ follow the reference signal $D_{ref}(t)$. The 10,000 ACs are divided into 10 clusters of 1000 ACs each. The bottom plot shows the global control signal.

The control in Fig. 10 aims to lower the demand to 75% in the reduction period and increase it to 95% during the recovery period. Note that the controller was designed with the model we calculated using (15) with the population parameters in Table 1, which assume a duty cycle close to 50%. Yet, when we use the resulting controller for a plant with very high duty cycle, the output remains close to the reference, showing the robustness of the proposed control design to this assumption.

Fig. 11 shows the power output of clusters 1 and 10 during the scenario presented in Fig. 9. Each curve is normalised to the maximum power of the cluster it represents. The demand peaks at the cluster level in Fig. 11 evidence that if the clusters had to be defined geographically (e.g. a group of adjacent city blocks), the presented algorithm may not be suitable, since high peaks might appear in the infrastructure powering that geographical area. Alternatively, the clusters can be defined so that each geographical area includes devices from many or all of the clusters. If such a cluster layout is possible (e.g. communications-wise), a peak at the cluster level should not have any negative impact on the distribution infrastructure.




Fig. 11. Output power demand D for two clusters (each normalised to maximum demand of cluster) under the LM scenario depicted in Fig. 9.

Another consideration when defining the clusters is the distribution of the AC parameters in each cluster. Our analysis assumes that the ACs are randomly chosen from the population to form the clusters and therefore each cluster shares the statistical properties of the population (e.g. same mean C). If a different clusterisation (e.g. some clusters with high and some other with low mean thermal capacitance) is required, additional analysis should be carried out in order to determine how this heterogeneity affects the results presented in this section.

6. A method to quantify the trade-off between discomfort and controllability

One advantage of using temperature set point offset as a control signal is that the entity implementing the LM (e.g. the utility) is able to obtain an estimate of the comfort impact straightforwardly (just by observing $u_i(t)$), since the temperature of any AC in the cluster i converges towards the hysteresis band $[\theta_- + u_i(t), \theta_+ + u_i(t)]$.

It is apparent, then, that changes in $u_i(t)$ will have an impact on the comfort of the occupants in the i th cluster. One of the ways of quantifying that impact is by the Predicted Percent Dissatisfied (PPD), a discomfort metric defined by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [26] that estimates the percentage of people that would vote that they are uncomfortably cold or hot if they were surveyed. The PPD is calculated using the Predicted Median Vote (PMV), which in turn is a function of temperature, clothing, humidity, activity level of the occupants, air velocity and other parameters [26].

In this section we propose a generalisation of the PPD that allows us to estimate the comfort range of the occupants of all the conditioned spaces in one cluster, regardless of the individual set point temperatures of each AC.

We assume that the occupants have chosen the set point they are most comfortable with. In other words, $u_i(t) = 0$ achieves the minimum dissatisfaction (PPD) possible.² Fig. 12 plots the range spanned by a set of PPD curves representing different combinations of clothing level, occupant's activity, air velocity and relative humidity according to the values shown in Table 3. All of the curves in Fig. 12 were superimposed, making their minima coincide at $u_i(t) = 0$, and the range that they collectively span was shaded, representing the range of comfort impact for a value of $u_i(t)$. In other words, by looking at the maximum and minimum values of $u_i(t)$ during a LM scenario, we can get an upper bound of how uncomfortable occupants get during that period.

For example, from Fig. 9, we see that $\max(u_1(t)) = 1$ and $\max(u_{10}(t)) = 0.5$. By looking at Fig. 12 one can conclude that in the worst case (assuming that the maximum value of $u_i(t)$ were maintained for long enough so that the temperature catches up with it), the PPD of the set point temperature would have risen from 5% to 6.37% for cluster 10 and to 10.15% for cluster 1.

Because u_1 and u_{10} are boundaries for any other individual signal, their comfort impacts bound that of any other cluster. An alternative way of calculating these boundaries without looking at the cluster signal is by “ceiling” the global control signal $u(t)$ to the closest multiple of $\Delta'u$ for the upper bound (i.e. $u_1(t)$) and “flooring” it for the lower bound (i.e. $u_{10}(t)$).

To the best of our knowledge, this is the first attempt to quantify the comfort impact when a population of ACs are controlled for LM by temperature set point offset changes. Applying the method presented in this section, a controller can be designed to regulate

² Note that the theoretical minimum PPD is 5%, modelling that it's impossible to keep everyone happy.




Fig. 12. Range of change in predicted percent dissatisfied when the temperature set point is changed for the values for clothing level, metabolic heat generation, air velocity and relative humidity presented in Table 3.

Table 3
Parameter range for calculating PPD.

Parameter	Value	Description
Clothing level	{0.36, 0.61} (clo)	0.36 = Walking shorts, short-sleeve shirt. 0.61 = Trousers, long-sleeve shirt
Metabolic heat generation	{1, 2} (met)	1 = reading/writing. 2 = mild housecleaning
Air velocity	{0, 0.2} (m/s)	
Relative humidity	{40, 60, 80} (%)	

the power demand to satisfy a given constraint on maximum comfort impact.

7. Conclusions

In this paper we have considered modelling and control of the aggregate power demand of a population of ACs. The ACs independently regulate temperature using relay-type actuators and a thermostat with a hysteresis band around a temperature set point.

Our model characterises the aggregate power consumption of the population of ACs by describing how the proportion of operating ACs varies over time, following a step change in the temperature set point of all the ACs. The ACs are assumed to be operating in steady state before the step change.

Under the additional assumption that the thermal capacitances are distributed log-normally in the population, we derived an explicit formula for the transient response to a common set point offset, a response which we further characterise as an underdamped oscillation. To the best of our knowledge, this is the first work that presents mathematical approximations for the period and amplitude envelope of such response. These approximations are shown to satisfactorily capture the dynamics of the response of a numerically simulated population of 10,000 ACs over a realistic range of parameter values. Moreover, we presented a way to derive the transfer function of an LTI second order system that characterises the aggregate power response of the population. Not only does this model provide rule-of-thumb information about the response (e.g. settling time) with no need for intensive numerical computations, but it is also apt for control design.

We used such a model to design an internal model control structure to compute common offsets to a temperature set point to be broadcast to the ACs as a control signal. We demonstrated that using this approach, the aggregate power output of the population of ACs can be satisfactorily controlled to provide a reduction

in the demand over a pre-specified period of time. The proposed model-based control displays robust performance even when the parameters characterising the population violate many of the simplifying modelling assumptions.

The proposed controller relies on fine resolution changes to the temperature set point of the ACs. In cases where the ACs only accept coarse set point offsets (for example 0.5 °C) we presented an implementation technique based on dividing the population into logical clusters and sending each cluster different control signals, each of them with admissible resolution. The aggregate power output does not show performance degradation compared to using a global, fine-grained, signal. Clustering, however, does incur in a comfort penalty that in a practical design needs to be deliberately traded off and can result in localised demand peaks if not appropriately distributed across the electricity network. Our method to quantify the range of discomfort for a given input signal enables the design of more advanced controllers that explicitly take into account the comfort impact associated with the control signal.

One of the lines we would like to explore in the future is the feasibility of using Corollary 2 to determine the parameters of a population from a system identification experiment (i.e., from the step response of a population of ACs, identify a second order system and from its transfer function determine the parameters describing the population).

The proposed approach provides a pathway to cost-effectively manage network demand, potentially deferring upgrades in the electricity transmission and distribution infrastructure. Moreover, this research is directly applicable to other types of TCLs such as fridges, cool rooms, water and space heaters.

Appendix A. Proofs and derivations

A.1. Proof of Proposition 1

Proof. From (5) and (6), the operational state $m_i(t)$ of the i th AC in the population is defined by

$$m_i(t) = \begin{cases} 0 & \text{if } 2k - 1 \leq x_i(t) < 2k, \\ 1 & \text{if } 2k \leq x_i(t) < 2k + 1, \end{cases} \quad (\text{A.1})$$

for $k = 1, 2, \dots$

In other words, we can determine whether or not the i th AC is turned on based on $x_i(t)$ being in an odd or even interval, if $x_i(t) \geq 1$. This is, however, not the case if $x_i(t) < 1$. When (6) is applied to every point in the temperature distributions after the step change shown in Fig. 1b, the distribution of the initial values of $x(t)$ for all of the ACs, namely $x(0)$, is as shown in Fig. A.13. Note that, initially, only one third of the ACs that satisfy $x_i(0) < 1$ are turned on. Therefore, under assumptions H.1–H.6, the probability that a randomly-picked AC is operating at time t given $x(t) < 1$ is

$$\Pr[m(t) = 1 | x(t) < 1] = \frac{\Pr[x(t) < 1]}{3}. \quad (\text{A.2})$$




Fig. A.13. Distribution of $x(0^+)$ (immediately after the temperature set point step change at $t = 0$).

Thus, for a number of ACs sufficiently large, the proportion $D(t)$ of operating ACs in the population at time t is equivalent to the probability of an AC being operating at time t ; namely,

$$D(t) = \frac{\Pr[x(t) < 1]}{3} + \sum_{k=1}^{\infty} \Pr[x(t) < 2k+1] - \Pr[x(t) < 2k].$$

A.2. Proof of Corollary 1

Proof. The fact that v is log-normally distributed follows from (4) and basic properties of the expected value and standard deviation of a log-normal distribution.

Let us now show the equality $\sigma_v/\mu_v = \sigma_C/\mu_C$ in (8). Considering that in (4) only the parameter C is distributed in the population (because of H.1, H.3, H.4), we have that

$$\frac{\sigma_v}{\mu_v} = \frac{s.d. \left[\frac{\theta_a - \theta_{ref}}{CR} \right]}{E \left[\frac{\theta_a - \theta_{ref}}{CR} \right]} = \frac{s.d. [1/C] \frac{\theta_a - \theta_{ref}}{R}}{E[1/C] \frac{\theta_a - \theta_{ref}}{R}} = \frac{s.d. [1/C]}{E[1/C]} \quad (A.3)$$

where $s.d.[\cdot]$ and $E[\cdot]$ are the standard deviation and expected value operators. Thus, since C is log-normally distributed,

$$\frac{s.d. [1/C]}{E[1/C]} = \frac{s.d. [C]}{E[C]} = \frac{\sigma_C}{\mu_C} = \sigma_{rel}. \quad (A.4)$$

From the fact that v has a log-normal distribution, it follows from (5) that $x(t)$ will be approximately log-normal for t large enough. Therefore, assuming that $x(t)$ is distributed log-normally, we have

$$\Pr[x(t) \leq y] = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left[\frac{\log(y) - \mu(t)}{\sqrt{2\sigma^2(t)}} \right] \quad (A.5)$$

where

$$\mu(t) = \log(\mu_x(t)) - \frac{1}{2} \log \left(1 + \frac{\sigma_x^2(t)}{\mu_x^2(t)} \right)$$

and

$$\sigma^2(t) = \log \left(1 + \frac{\sigma_x^2(t)}{\mu_x^2(t)} \right).$$

That is, $\mu(t)$ and $\sigma(t)$ are the mean and standard deviation of the normal distribution associated to the log-normal distribution characterising x .

According to (5), the mean $\mu_x(t)$ and variance $\sigma_x^2(t)$ of $x(t)$ evolve as follows:

$$\mu_x(t) = E[x(t)] = E[x(0)] + E[v]t = \mu_x(0) + \mu_v t, \quad (A.6)$$

$$\begin{aligned} \sigma_x^2(t) &= (s.d.[x(t)])^2 = (s.d.[x(0)])^2 + (s.d.[v])^2 t^2 \\ &= \sigma_x^2(0) + \sigma_v^2 t^2, \end{aligned} \quad (A.7)$$

Thus, we can approximate, for t large enough,

$$\frac{\sigma_x^2(t)}{\mu_x^2(t)} \approx \frac{\sigma_v^2}{\mu_v^2} = \sigma_{rel}^2. \quad (A.8)$$

and then, using (A.6) and (A.8) in (A.5) we obtain

$$\Pr[x(t) \leq y] \approx \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left[\frac{\log(y) - \log(\mu_x(0) + \mu_v t) + \frac{1}{2} \log(1 + \sigma_{rel}^2)}{\sqrt{2 \log(1 + \sigma_{rel}^2)}} \right]. \quad (A.9)$$

We can approximate $\log(1 + \sigma_{rel}^2) \approx \sigma_{rel}^2$ with an absolute error of the order of $(\sigma_{rel})^8/2$, arriving to

$$\Pr[x(t) \leq y] \approx \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left[\frac{\log(y) - \log(\mu_x(0) + \mu_v t)}{\sqrt{2\sigma_{rel}}} \right] \quad (A.10)$$

where we have disregarded the term $\frac{1}{2} \log(1 + \sigma_{rel}^2)$ from the numerator inside the erf function in (A.9) because it is negligible as compared to $\log(\mu_x(0) + \mu_v t)$.

Replacing (7) with (A.10) for $y = 1$, $y = 2k$ and $y = 2k + 1$, we obtain (9).

A.3. Proof of Proposition 2

Proof. The mean $\mu_x(t)$ and variance $\sigma_x^2(t)$ increase monotonically with time and, in fact, after a short initial period of time, only one value of k will account for the dominant contribution to the summation term in (9). Thus, we approximate the probability of randomly-chosen AC satisfying $k < x(t) < k + 1$ as

$$\begin{aligned} \Pr[k < x(t) < k + 1] &= \Pr[x(t) < k + 1] - \Pr[x(t) < k] \\ &\approx \frac{\operatorname{erf} \left(\frac{\log(k+1) - \tau}{\sqrt{2\sigma_{rel}}} \right) - \operatorname{erf} \left(\frac{\log(k) - \tau}{\sqrt{2\sigma_{rel}}} \right)}{2}, \end{aligned} \quad (A.11)$$

where

$$\tau = \log(\mu_x(0) + \mu_v t). \quad (A.12)$$

For even values of k , there will be a positive peak of power when t is such that (A.11) is a local maximum. The peak will have amplitude $D_k / \sum_{i=1}^n D_i \approx \Pr[k < x(t) < k + 1]$. Analogously, when k is odd, a negative peak of power will be encountered when t is such that (A.11) is maximum and its amplitude will be $D_k / \sum_{i=1}^n D_i \approx 1 - \Pr[k < x(t) < k + 1]$.

Therefore, for both even and odd values of k we need to find the value of τ for which (A.11) is maximised in order to compute the amplitude of the peaks. For that value of τ , it should be verified that

$$\frac{d}{d\tau} \left[\frac{\operatorname{erf} \left(\frac{\log(k+1) - \tau}{\sqrt{2\sigma_{rel}}} \right) - \operatorname{erf} \left(\frac{\log(k) - \tau}{\sqrt{2\sigma_{rel}}} \right)}{2} \right] = 0$$

from which we obtain

$$\begin{aligned} \left(\frac{\log(k+1) - \tau}{\sqrt{2\sigma_{rel}}} \right)^2 &= \left(\frac{\log(k) - \tau}{\sqrt{2\sigma_{rel}}} \right)^2 \\ \text{and therefore} \\ \frac{\log(k+1) - \tau}{\sqrt{2\sigma_{rel}}} &= \pm \frac{\log(k) - \tau}{\sqrt{2\sigma_{rel}}}. \end{aligned} \quad (A.13)$$

The equality (A.13) can only be satisfied when its right hand side is negative, which yields the value of τ for the k th peak in the output:

$$\tau^*(k) = \frac{\log(k+1) + \log(k)}{2}. \quad (A.14)$$

Replacing (A.14) in (A.11) and operating we obtain

$$\begin{aligned} \max_t \Pr[k < x(t) < k + 1] &\approx \operatorname{erf} \left(\frac{\log(k+1) - \log(k)}{2\sqrt{2\sigma_{rel}}} \right) \\ &= \operatorname{erf} \left(\frac{\log(1 + 1/k)}{2\sqrt{2\sigma_{rel}}} \right). \end{aligned}$$

For values of $k > 1$, we can approximate $\log(1 + 1/k) \approx 1/k$ with an absolute error of the order of $1/(2k^2)$, and then

$$\max_t \Pr[k < x(t) < k + 1] \approx \operatorname{erf} \left(\frac{1}{2k\sqrt{2\sigma_{rel}}} \right). \quad (A.15)$$

On the other hand, from (A.12) and (A.14) we know that

$$\tau^*(k) = \log(\mu_x(0) + \mu_v t^*(k)) = \frac{\log(k+1) + \log(k)}{2} = \log(\sqrt{k^2 + k}),$$

from which we obtain

$$\mu_x(0) + \mu_v t^*(k) = \sqrt{k^2 + k}, \quad (\text{A.16})$$

where $t^*(k)$ is the time of the k -th peak.

For $k > 1$ we can approximate $\sqrt{k^2 + k} \approx k + \frac{1}{2}$ with an absolute error of the order of $1/(8k)$ using the first two terms of the Taylor series expansion. Thus, from (A.15) and (A.16) we obtain

$$\begin{aligned} \max_t \Pr[k < x(t) < k+1] &\approx \text{erf}\left(\frac{1}{2k\sqrt{2}\sigma_{rel}}\right) \\ &\approx \text{erf}\left(\frac{1}{2(\mu_x(0) + \mu_v t^*(k) - \frac{1}{2})\sqrt{2}\sigma_{rel}}\right). \end{aligned} \quad (\text{A.17})$$

This last equation says that the k th peak of power occurs when $\mu_x(0) + \mu_v t - 1/2 = k$, which is to say that $t^*(k) = (k + 1/2 - \mu_x(0))/\mu_v$. Therefore, the period of the oscillations is given by $t^*(k+2) - t^*(k) = 2/\mu_v$, which shows (10).

Finally, inequality (11) is shown by evaluating the amplitude of the k th peak, which is

$$E\left[D_k / \sum_{i=1}^n D_i\right] \approx \text{erf}\left(\frac{1}{2k\sqrt{2}\sigma_{rel}}\right), \text{ if } k \text{ is even} \quad (\text{A.18})$$

or

$$E\left[D_k / \sum_{i=1}^n D_i\right] \approx 1 - \text{erf}\left(\frac{1}{2k\sqrt{2}\sigma_{rel}}\right), \text{ if } k \text{ is odd.} \quad (\text{A.19})$$

Taking into account that $\text{erf}(1/z)$ decreases monotonically for positive values of z , expressions (A.18) and (A.19) show how the amplitude of successive peaks decreases with k . Moreover, the larger σ_{rel} , the faster it decreases. \square

A.4. Proof of Corollary 2

For a given value of $z = z_1$, the upper part of the envelope of the oscillatory damped step response is $\text{erf}(1/z_1)$ according to (11). After one semi-period $1/\mu_v$ (see (10)), we have $z_2 = z_1(t + 1/\mu_v)$ and, applying (12), $z_2 = z_1 + 2\sqrt{2}\sigma_{rel}$. Thus, the ratio between the distance of two successive peaks to the steady state value (0.5), namely D_{k+1} and D_k is:

$$D_{k+1}/D_k = \frac{|\text{erf}\left(\frac{1}{z_1 + 2\sqrt{2}\sigma_{rel}}\right) - 0.5|}{|\text{erf}\left(\frac{1}{z_1}\right) - 0.5|}. \quad (\text{A.20})$$

As seen in Fig. 4, the approximation of the step response with a second order model has a reasonable fit for values of $z \in [0.6, 1.4]$. Thus, we propose a mid-interval value $z_1 = 0.9$ and replace it in (A.20) to obtain

$$D_{k+1}/D_k = \frac{|\text{erf}\left(\frac{1}{0.9 + 2\sqrt{2}\sigma_{rel}}\right) - 0.5|}{|\text{erf}\left(\frac{1}{0.9}\right) - 0.5|} \triangleq r(\sigma_{rel}).$$

Then, we can compute the damping factor

$$r = \exp\left[\frac{-\pi\zeta}{\sqrt{1-\zeta^2}}\right]$$

which yields

$$\zeta^2 = \frac{\log^2(r)}{\pi^2 + \log^2(r)}.$$

Since (10) states that the period can be approximated as $2/\mu_v$, the undamped natural frequency can now be computed as

$$\omega_n = \frac{\omega_d}{\sqrt{1-\zeta^2}}$$

where

$$\omega_d = \pi\mu_v.$$

With that, we have the denominator $s^2 + 2\zeta\omega_n s + \omega_n^2$ of the transfer function as parametrised by Corollary 2.

Let us compute the coefficients b_0, b_1 and b_2 of the numerator. We know that the value of $D(t)$ when the system is in steady state is the expected value of the stochastic expression $T_{on}/(T_{on} + T_{off})$, where T_{on}^i and T_{off}^i are the times that it takes the i th AC in the population to go from one end of the hysteresis band to the other when the device is operating (T_{on}^i) or not (T_{off}^i). Since, T_{on} and T_{off} depend on the reference temperature θ_{ref} , we can write them as a function of θ_{ref} [24]. Thus

$$T_{on}(\theta_{ref}) = CR \log\left(\frac{PR + \theta_{ref} + H/2 - \theta_a}{PR + \theta_{ref} - H/2 - \theta_a}\right) \quad (\text{A.21})$$

and

$$T_{off}(\theta_{ref}) = CR \log\left(\frac{\theta_a - \theta_{ref} + H/2}{\theta_a - \theta_{ref} - H/2}\right). \quad (\text{A.22})$$

Then, replacing these expressions in

$$D_{ss}(\theta_{ref}) = T_{on}(\theta_{ref}) / (T_{on}(\theta_{ref}) + T_{off}(\theta_{ref})). \quad (\text{A.23})$$

we obtain (16).

Let $Y(s) = G_p(s)0.5/s$ in (13) and $\mathcal{L}^{-1}\{Y(s)\} = y(t)$. Then, from the initial value theorem [11], we have

$$D_{ss}(\theta_{ref})/2 = y(0^+) = \lim_{s \rightarrow \infty} 0.5G_p(s) = 0.5b_2, \quad (\text{A.24})$$

therefore $b_2 = D_{ss}(\theta_{ref})$.

On the other hand, from the final value theorem [11], we have

$$D_{ss}(\theta_{ref} + 0.5) - D_{ss}(\theta_{ref}) = y_\infty = \lim_{s \rightarrow 0} 0.5G_p(s) = 0.5b_0/\omega_n^2, \quad (\text{A.25})$$

therefore $b_0 = (D_{ss}(\theta_{ref} + 0.5) - D_{ss}(\theta_{ref}))\omega_n^2/0.5$.

Combining the initial value theorem and the Laplace transform of a derivative, we have

$$\dot{y}(0^+) = \lim_{s \rightarrow \infty} 0.5G_p(s) - 0.5b_2 = 0.5(b_1 - 2b_2\zeta\omega_n), \quad (\text{A.26})$$

which gives us $b_1 = \dot{y}(0^+)/0.5 + 2b_2\zeta\omega_n$.

Looking at Fig. A.13, we can see that 25% of the ACs are operating at $t = 0^+$. Since we know that the distribution of $x(t)$ moves to the right at a mean speed μ_v , we can approximate $\dot{y}(0^+) \approx \frac{y(0^+ + \Delta t) - y(0^+)}{\Delta t} = -\frac{(0.25 - 0.25\mu_v\Delta t) - 0.25}{\Delta t} = 0.25\mu_v$. Thus, $b_1 = 0.5\mu_v + 2b_2\zeta\omega_n$.

References

- [1] World Energy Outlook 2006, International Energy Agency, Organisation for Economic Co-operation and Development, Paris, France; 2006.
- [2] Kirby BJ. Spinning reserve from responsive loads. Tech. Rep. ORNL/TM-2003/19. Oak Ridge National Laboratory; 2003.
- [3] Callaway D, Hiskens I. Achieving controllability of electric loads. Proc IEEE 2011;99(1):184–99.
- [4] Callaway DS. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. Energy Convers Manage 2009;50(5):1389–400.
- [5] Lu N, Katipamula S. Control strategies of thermostatically controlled appliances in a competitive electricity market. In: Power engineering society general meeting, 2005, vol. 1. IEEE 2005. p. 202–7.
- [6] Heffner G, Goldman C, Kirby B, Kintner-Meyer M. Loads providing ancillary services: review of international experience. Lawrence Berkeley National Laboratory.
- [7] Schisler K, Sick T, Brief K. The role of demand response in ancillary services markets. In: Transmission and distribution conference and exposition, 2008. IEEE/PES; 2008. p. 1–3.

- [8] Xu Z, Ostergaard J, Tøgeby M. Demand as frequency controlled reserve. *IEEE Trans Power Syst* 2010(99):1–10.
- [9] Ruiz N, Cobelo I, Oyarzabal J. A direct load control model for virtual power plant management. *IEEE Trans Power Syst* 2009;24(2):959–66.
- [10] Ihara S, Schweppe F. Physically based modeling of cold load pickup. *IEEE Trans Power Apparatus Syst* PAS-100 1981(9):4142–50.
- [11] Goodwin GC, Graebe SF, Salgado ME. *Control system design*. 1st ed. Upper Saddle River, NJ, USA: Prentice Hall PTR; 2000.
- [12] Chong C-Y, Debs AS. Statistical synthesis of power system functional load models. In: *Proceedings of the 18th IEEE conference on decision and control*, vol. 18; 1979. p. 264–9.
- [13] Malhamé RP. A statistical approach for modeling a class of power system loads. Ph.D. thesis, Georgia Institute of Technology; 1983.
- [14] Malhamé RP, Chong C-Y. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Trans Autom Control* 1985;30(9):854–60.
- [15] Chong C-Y, Malhamé RP. Statistical synthesis of physically based load models with applications to cold load pickup. *IEEE Trans Power Apparatus Syst* PAS-103 1984(7):1621–8. doi:[10.1109/TPAS.1984.318643](https://doi.org/10.1109/TPAS.1984.318643).
- [16] Bect J, Baili H, Fleury G. Generalized Fokker-Planck equation for piecewise-diffusion processes with boundary hitting resets. In: *Proceedings of the international symposium on mathematical theory of networks and systems (MTNS 2006)*; 2006. p. Paper WeA05.6.
- [17] Molina-García A, Kessler M, Fuentes J, Gómez-Lázaro E. Probabilistic characterization of thermostatically controlled loads to model the impact of demand response programs. *IEEE Trans Power Syst* 2011;26(1):241–51.
- [18] El-Ferik S, Malhamé R. Identification of alternating renewal electric load models from energy measurements. *IEEE Trans Autom Control* 1994;39(6):1184–96. doi:[10.1109/9.293178](https://doi.org/10.1109/9.293178).
- [19] El-Ferik S, Hussain S, Al-Sunni F. Identification of physically based models of residential air-conditioners for direct load control management. In: *Proceedings of the 5th Asian control conference*, vol. 3; 2004. p. 2079–87.
- [20] Burke W, Auslander D. Robust control of residential demand response network with low bandwidth input. In: *Proceedings of the ASME dynamic systems and control conference*; 2008. p. 1187–9.
- [21] Laurent J-C, Desaulniers G, Malhamé RP, Soumis F. A column generation method for optimal load management via control of electric water heaters. *IEEE Trans Power Syst* 1995;10(3):1389–400.
- [22] Ljung L. *System identification: theory for the user*. Upper Saddle River, NJ, USA: Prentice-Hall; 1986.
- [23] Bashash S, Fathy HK. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In: *Proc. American control conference*; 2011.
- [24] Kundu S, Sinityn N, Backhaus S, Hiskens I. Modelling and control of thermostatically controlled loads. In: *Proc. power systems computation conference*; 2011.
- [25] Koch F, Mathieu JL, Callaway DS. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In: *Proc. power systems computation conference*; 2011.
- [26] Parsons R, editor. 2005 ASHRAE handbook: fundamentals. American Society of Heating, Refrigerating and Air-Conditioning Engineers; 2005.
- [27] Bergero F, Kofman E. PowerDEVS. A tool for hybrid system modeling and real time simulation. *Simul: Trans Soc Model Simul Int* 2011;87(1–2):113–32.
- [28] Franklin GF, Powell DJ, Emami-Naeini A. *Feedback control of dynamic systems*. 4th ed. Upper Saddle River, NJ, USA: Prentice Hall PTR; 2001.