A sensitivity analysis of the dynamics of a population of thermostatically-controlled loads

Cristian Perfumo
CSIRO Energy Technology
Australia
cristian.perfumo@csiro.au

Julio Braslavsky
CSIRO Energy Technology and
University of Newcastle
julio.braslavsky@csiro.au

John K. Ward
CSIRO Energy Technology
Australia
john.k.ward@csiro.au

***Abstract*—Load control of populations of thermostatically-controlled loads (TCLs) is considered a promising approach to match generation and consumption in electricity grids from the demand side. However, when these loads become synchronised they present a decaying oscillatory aggregate demand, which results in undesired power peaks. In this paper we describe the nature of these oscillations and develop a list of factors that determine its shape. We perform a sensitivity analysis which allows us to identify important relations between the physical parameters of the TCLs and their aggregate dynamics. Beyond describing fundamental behaviour, these relations can help develop and validate analytical expressions that facilitate control design, enabling the use of TCLs for demand response.**

I. INTRODUCTION

The widespread use and inherent energy storage capabilities of thermostatically-controlled loads (TCLs) such as air conditioners (ACs) and fridges, make these loads promising candidates for demand response [1]. Controlling large groups of TCLs can help achieve the necessary balance between generation and consumption in electricity grids. This demand-side approach can facilitate renewable energy integration and defer expensive infrastructure upgrades [2].

However, remotely controlling groups of TCLs may result in large transients in the collective power demand, which may lead to undesired power peaks. Figure 1, shows simulated aggregate demand responses of a population of (10,000) TCLs in three of these likely unwanted situations, namely

- when, following an extended power outage, the power of all of the ACs is restored (Figure 1(a)),
- when all of the devices are subject to a step change in the temperature set point (Figure 1(b)), and
- when a load control event is finished and full independent control is returned to the ACs while they are still substantially synchronised (Figure 1(c)), which could be due to bad planning of the DLC event or a communications problem terminating the event prematurely.

We refer to these events as *synchronisation events* because they make a significant proportion of the TCLs in the population synchronise (i.e., turn on or off at the same time), causing an observable change in their aggregate demand. Note in Figure 1 that this observable change (sometimes referred to as “cold load pickup”) comes in the form of damped oscillatory transients, which eventually converge the *steady state aggregate demand* of the population. Furthermore, similar oscillations may occur with loads under randomised

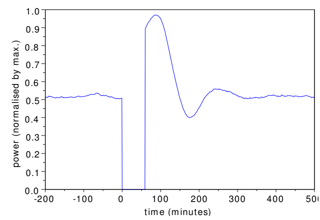

(a) Power outage

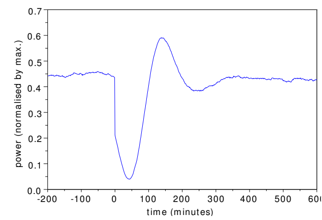

(b) Common step change in temperature set point

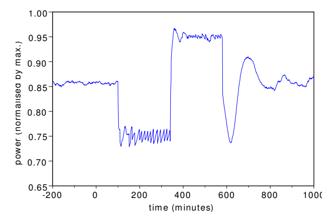

(c) Bad load control event

Fig. 1. Examples of synchronisation events that can cause oscillations in the aggregate power demand of a population of TCLs.

(but uncoordinated) autonomous control, as illustrated in [1] for populations of plug-in electric vehicles.

Developing models of this oscillatory response and the dynamics in general of the aggregate power of a population of TCLs is currently an area of significant interest, as these models prove useful not only for simulation but also control design [2]–[8]. Model-based controllers show great performance, enabling the use of populations of TCLs to shift load from peak to off-peak times and to follow the intermittent output of renewable generation.

This paper presents a comprehensive sensitivity analysis of the basic dynamic characteristics of the aggregate response of a population of ACs to a synchronisation event. In particular, we focus on identifying the factors that affect most significantly the shape of this response. Our approach in this study is mostly numerical: we simulate populations of ACs and observe the changes in the aggregate response to variations in the simulation parameters.

In particular, we analyse how the response is influenced by the type and level of heterogeneity in the population, the ratio between mean thermal parameters, the magnitude of the synchronisation event and the steady-state aggregate demand. Our main three findings are that a) more heterogeneous populations present more damped oscillatory responses, b) heterogeneity does not affect the period of the oscillations, and c) the type of distribution of the parameters R, P and C in the population does not affect the oscillations significantly.

II. Simulation methodology

We consider a population of n ACs regulating room temperatures by hysteretic control with a thermostat and relay actuator. The relay of the i-th AC, $i \in \{1, 2, \ldots, n\}$, is represented by the discrete state variable $m_i \in \{0, 1\}$, which switches *on* the AC compressor ($m_i = 1$) or *off* ($m_i = 0$) to maintain the temperature θ_i within the pre-specified hysteresis band $[\theta_i^-, \theta_i^+]$. This temperature band is centred on the nominal temperature set-point $\theta_i^{\mathrm{r}} = (\theta_i^- + \theta_i^+)/2$. Such dynamics can be described by the well-known hybrid state model (e.g., [9], [10])

$$\frac{d\theta_i(t)}{dt} = -\frac{1}{C_i R_i}\big[\theta_i(t) - \theta_a(t) + m_i(t)R_i P_i - w_i(t)\big], \quad (1)$$

$$m_i(t^+) = \begin{cases} 0, & \text{if } \theta_i(t) \le \theta_i^- + u(t), \\ 1, & \text{if } \theta(t) \ge \theta_i^+ + u(t), \\ m_i(t), & \text{otherwise,} \end{cases} \quad (2)$$

where θ_a is the ambient temperature outside the rooms (°C) (assumed common to all ACs), 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 input signal w_i represents unpredictable thermal disturbances (heat gains or losses). The control signal u (common to all ACs) introduces small temporary temperature set-point offsets to the population during LC events (using this type of control signal has been proposed in a number of recent works [2], [3], [7]).

We simulate individually 10000 ACs as an array of instances of the continuous and discrete state dynamic equations (1) and (2) using the event-based simulation environment PowerDEVS [11]. The thermal parameters R, P and C in (1) are assumed to be lognormally distributed in the population (as done in [2]). The relative standard deviation σ_{r} of these three parameters with respect to their mean is assumed to be the same (as done in [2], [4]). We use the parameters in Table I for our simulations, which were adapted from [2] to the characteristics of Australian suburban houses [12].

The aggregate electrical power demand (kW) of the population of ACs as a fraction of the maximum electrical power demand (when all the AC are on) is given by

$$D^{\mathrm{ac}}(t) = \frac{\sum_{i=1}^{n} m_i(t)\frac{P_i}{\mathrm{COP}_i}}{\sum_{i=1}^{n} \frac{P_i}{\mathrm{COP}_i}}, \quad (3)$$

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

III. The nature of the oscillations in the aggregate power demand of a population of ACs

Let us analyse the oscillations for the step case in Figure 1(b) (those observed in the power outage and bad-DLC scenarios shown in Figures 1(a) and 1(c) are of similar nature). Figure 2 presents snapshots of the temperature distributions in the population at various stages during the transients. Each subfigure 2(a)–2(a) contains three subplots. The left subplot shows, the aggregate demand response (normalised by the maximum demand), and the time t at which the histograms are plotted. The middle subplot shows the temperature histogram of the ACs that are on at the corresponding time t, and the right subplot shows the histogram of the ones that are off. To aid visualisation, all of the ACs in Figure 2 are assumed to have the same temperature set point and hysteresis band (however, this is no the case in the rest of the paper, where both parameters are distributed in the population). The step raises the set point by 0.5 ^oC, half the width of this hysteresis band.

Just before the step takes place, the temperatures are distributed almost uniformly within the hysteresis band $[19.5; 20.5]$, as can be seen in Figure 2(a). Immediately after the step (Figure 2(b)), the "cooler" half of the ACs that were on turned off, since their temperatures are now below the newly-defined hysteresis band. This causes an instantaneous reduction in power demand at the time of the step.

After the step, there is a period in which no off ACs turn on, but the on ACs that are close to the new lower hysteresis boundary (20 ^oC) continue to turn off. This causes the demand to keep decreasing for some time after the step.

When the off devices start to reach the newly defined upper boundary of the hysteresis width (21 ^oC), they turn on and cause the aggregate demand to decrease more slowly until it reaches a trough (Figure 2(c)). After this local minimum there are, on average, more ACs turning on than turning off, which causes the demand to rise (Figures 2(d) and 2(e)) until it reaches a peak (Figure 2(f)).

Eventually, as shown in Figure 2(h), the histograms (and the aggregate demand) are back to steady state, now half a degree higher, as intended with the step.

Summarising, when a significant portion of a heterogeneous

Param.	Value	Description
R	2 ^oC/kW	Mean thermal resistance (distributed lognormally according to σ_{r})
C	3.6 kWh/^oC	Mean thermal capacitance (distributed lognormally according to σ_{r})
P	6 kW	Mean thermal power (distributed lognormally according to σ_{r})
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	26 ^oC	Outside temperature
σ_w	0.01 $^oCs^{-\frac{1}{2}}$	Standard deviation of the noise process w in Eq. (1)
σ_{r}	0.2	Standard deviation of log-normal distributions as a fraction of the mean value for R, C and P
COP	2.5	Coefficient of performance (thermal power on electrical power)
n	10000	Number of ACs in the population

TABLE I
Simulation parameters.

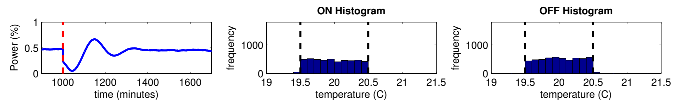

(a) $t = 1000$

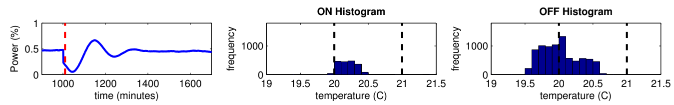

(b) $t = 1010$

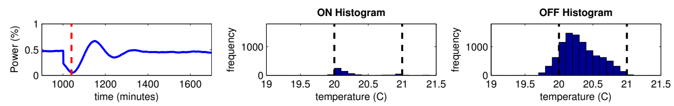

(c) $t = 1040$

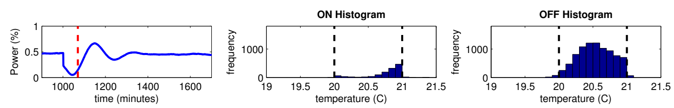

(d) $t = 1070$

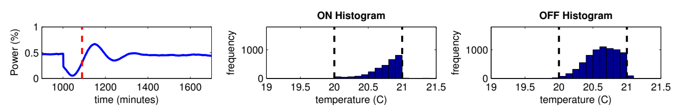

(e) $t = 1090$

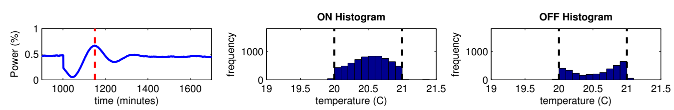

(f) $t = 1150$

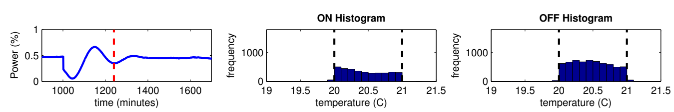

(g) $t = 1250$

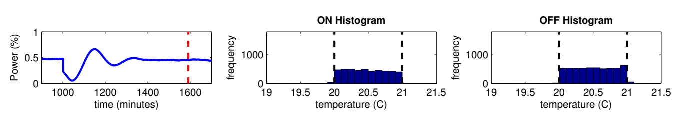

(h) $t = 1590$

Fig. 2. Temperature histograms of 10000 ACs after a 0.5 ^oC step increment in their temperature set point at $t = 1000$. Subfigure rows (a) to (h) show on the left the aggregate demand response with a red dashed line indicating the snapshot time for the histograms shown on the centre and right columns: temperature distribution for ACs that are "on" (centre) and "off" (right).

group of ACs is forced to switch on or off at the same time, the equilibrium between the devices switching on and those switching off is disrupted. When more ACs switch on than off, the aggregate demand increases, and when more devices turn off than on, it decreases. The heterogeneity in the population causes these imbalances to decay over time, until a new equilibrium is reached and the aggregate demand stops oscillating (i.e., the system reaches a dynamic steady state). Following, we analyse a number of factors that affect the particular shape of these oscillations.

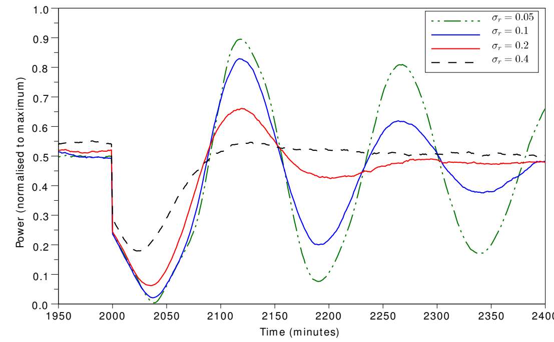

Fig. 3. Effect on the aggregated response when varying σ_{r}, the standard deviation of the parameters R, C and P as a fraction of their means.

IV. Factors that influence the characteristics of the oscillations in aggregate power demand

As we will show next, the transient aggregate demand response of a population of ACs to a common step change in temperature set point depends on a number of factors that affect its shape, namely

1) *level of heterogeneity in the population*,
2) *type of heterogeneity in the population*,
3) *ratio between mean thermal parameters*,
4) *magnitude of the synchronisation event* and
5) *steady-state aggregate demand.*

In this section we perform a sensitivity analysis of the impact of each of these factors on the aggregate power demand of a simulated population of ACs. In our analysis we will use a common step change in the temperature set point (such as the one in Figure 1(b)) as the example synchronisation event because the aggregate power of a population of ACs and other TCLs may be effectively controlled by manipulating a common temperature set point offset [2]–[4], [13]. However, most of the conclusions drawn from the response to a step change are also applicable to other synchronisation events such as power outages (unless otherwise stated).

A. Level of heterogeneity in the population

In any realistic population of TCLs there is some degree of heterogeneity (e.g., different ACs and dwelling characteristics). Even for "homogeneous" populations (i.e., identical devices conditioning identical spaces) [14], internal heat gains such as people, electrical appliances account for some heterogeneity (hence the stochastic variable $w(t)$ in (1)).

The level of heterogeneity in a population directly influences for how long the ACs remain coherent upon synchronisation. Figure 3 shows the step response of a population of ACs, where their parameters P, C, and R are sampled from lognormal probability distributions. The relative standard deviation σ_{r} of these three parameters with respect to their mean is varied to illustrate the effect of heterogeneity.

After a synchronisation event, more heterogeneous populations (larger σ_{r}) lose coherence faster than less heterogeneous ones (in dynamics systems terminology, we say that their responses are more *damped*). This loss of coherence is due to the fact that the more alike the ACs are, the longer temperature

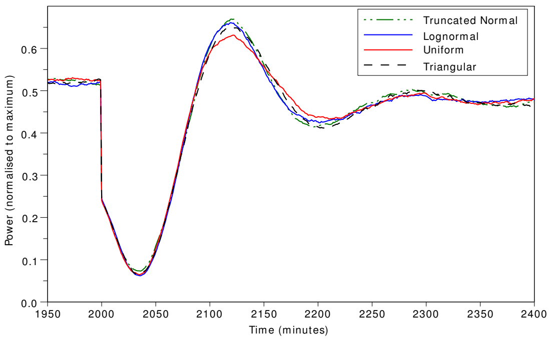

Fig. 4. Step responses of populations with thermal parameters distributed uniformly, lognormally, normally (truncated) and triangular (symmetric). The parameter means and standard deviations are as detailed in Table I.

dynamics will remain synchronised moving together from the on to the off state and back (in the extreme, unrealistic case of total homogeneity, the ACs will remain synchronised indefinitely). On the other hand, highly heterogeneous populations will rapidly lose coherence and reach steady-state power demand.

The level of heterogeneity does not seem to affect significantly the *frequency* of the oscillations. However, more heterogeneity causes an earlier first peak in the response, as the distribution of the on and off state transitions reach a detached balanced faster due to faster spreading of the sharp front of the right-moving off-state distribution.

B. Type of heterogeneity in the population

The heterogeneity of a population can be characterised by the distribution of the parameters that govern the behaviour of the devices (i.e., those in (1)). This characterisation is done using the lognormal distribution in [2], the uniform distribution in [15] and the normal distribution in [4].

We consider the effect of using different distributions while maintaining all other model properties (including mean and variance of these parameter values). Figure 4 compares the step responses when the distribution of R, P and C in the population is lognormal, uniform, normal (truncated) or triangular (symmetric).

As seen in Figure 4, the responses associated with normal, lognormal and triangular distributions are almost identical, whereas the uniform differs in that it exhibits a slightly lower local maximum after the step and the rate of demand decrease after such maximum is somewhat lower.

An explanation for this similarity lies in (1) where we can see that the parameters P, C and R appear grouped as (CR) and (RP). Figure 5(a) depicts the histograms of the product CR for the lognormal and uniform cases in Figure 4. We can see in Figure 5(a) that the product distributions are significantly similar considering how different the lognormal and uniform distributions are[1]. Similarly, Figure 5(b) shows that the distribution of the cycling times of both populations does not present significant differences either. The cycling time

[1]The general expression for the distribution of the product of two random variables is well-known but difficult to calculate [16], except for the case of two lognormally distributed variables (whose product is lognormal).

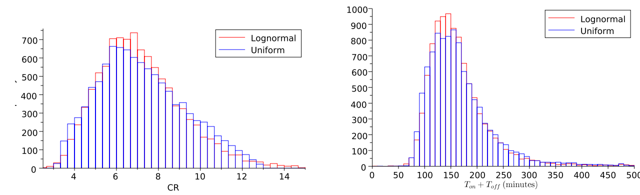

Fig. 5. Histograms of: (left) the parameter product CR when R and C have lognormal or uniform distributions and (right) the cycle times of the ACs, computed from (5) and (6) as $T_{\text{on}} + T_{\text{off}}$.

of each AC is computed as $T_{\text{on}} + T_{\text{off}}$, where T_{on} is the time taken by the AC to lower the temperature from θ_+ to θ_-, and T_{off} is the time for the temperature to raise from θ_- to θ_+ when the AC is off. We provide expressions for T_{on} and T_{off} later in the paper (Equations (5) and (6)).

Figures 5(a) and 5(b) indicate that the type of distribution of the individual parameters R, P and C does not have a substantial impact in the shape of the aggregate response.

C. Ratio between mean thermal parameters

Figure 6(a) shows how the aggregate response to a common 0.5 ^oC step increase in the temperature set point changes for populations with different mean values of P (everything else remaining the same). We see that larger mean values of P lead to lower steady state demand. To explain this, let us define the duty cycle of an AC as

$$D_c = T_{\text{on}}/(T_{\text{on}} + T_{\text{off}}). \tag{4}$$

where T_{on} and T_{off} are, again, the times required for the temperature to traverse the hysteresis band when the compressor is on or off respectively. The value of T_{on} can be computed assuming $w(t) = 0$ and solving (1) for initial conditions $\theta(0) = \theta_+$ and $m(0) = 1$ (the AC is on) [3], which yields

$$T_{\text{on}} = CR \log\left(\frac{PR + \theta_+ - \theta_a}{PR + \theta_- - \theta_a}\right). \tag{5}$$

Similarly, we can write T_{off} as

$$T_{\text{off}} = CR \log\left(\frac{\theta_a - \theta_-}{\theta_a - \theta_+}\right). \tag{6}$$

In (1) we can see that when the AC is on, larger P gives faster rate of temperature change, which implies smaller T_{on}. It takes little calculation to show from (5) that T_{on} decreases as P increases, which is seen by differentiating T_{on} with respect to P, which gives $dT_{\text{on}}/dP < 0$. On the other hand, it is clear from (6) that P has no effect on T_{off}. Thus, the duty cycle (4) of an AC decreases for larger values of P. Generalising, larger *mean* values of P in the population cause the *mean* duty cycle to decrease, which in turn implies a lower steady-state demand. In the particular case where all of the ACs have the same electrical power (i.e., P divided by the coefficient of performance), the mean duty cycle and normalised steady state demand coincide.

A second observation from Figure 6(a) is that larger values of P result in shorter periods in the oscillations of the aggregate response. As shown in Figure 2, the period of these

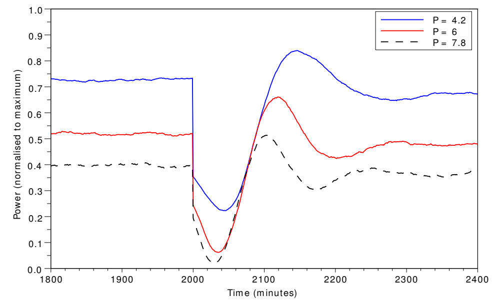

(a) Impact of mean P on the step response.

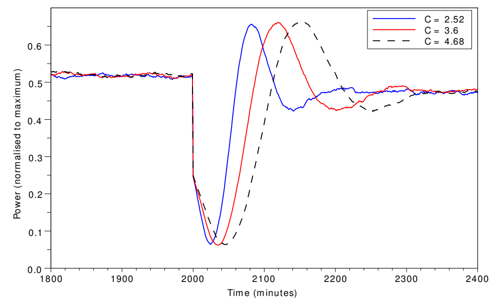

(b) Impact of mean C on the step response.

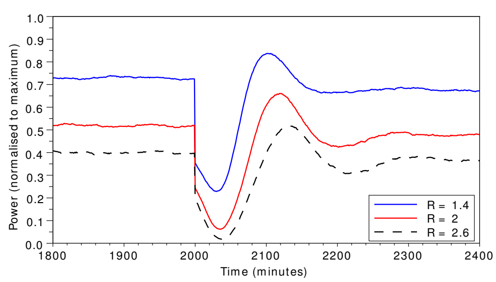

(c) Impact of mean R on the step response.

Fig. 6. Everything else remaining the same, impact of varying the mean values of P, C and R in the population (6(a), 6(b) and 6(c) respectively).

oscillations depend on the speed at which the temperatures move from one boundary of the hysteresis band to the other. A larger mean P results in temperatures going from θ_+ to θ_- in less time, thus reducing the period of the oscillations.

Regarding the influence of the mean value of C in the response, in (1) the temperature change rate (both when the AC is on and off) is inversely proportional to C. Thus, higher mean values of C make the ACs take longer to traverse the hysteresis band (i.e., slower oscillations). Figure 6(b) illustrates this: higher mean values of C in the population produce responses with slower frequency (with no effect in steady-state aggregate demand or magnitude of the peaks). Interestingly, in [2] it is proven that for homogeneous populations, larger CR implies slower convergence. The observation just made from Figure 6(b) suggests that this theoretical result could be generalised for heterogeneity (in Figure 6(b), larger CR follows from C increasing and R staying the same).

With respect to the average thermal resistance in the population, larger mean values of R result in lower frequency

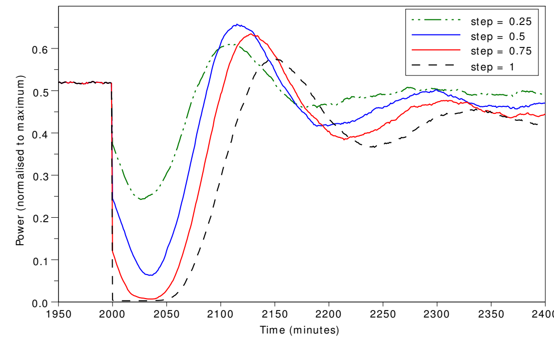

Fig. 7. Effect of the amplitude of a step input (in ^oC) in $u(t)$.

in the aggregated response. Much like in the case of C, the reason for this can be seen in (1): larger R implies slower temperature changes in absolute value. In fact, the aforementioned theoretical result in [2] can also be interpreted as: larger R and equal C implies slower convergence. This effect is illustrated in Figure 6(c). We also see in Figure 6(c) that higher values of R yield lower steady state power demand. This is because higher values of R represent better insulated spaces, saving energy by reducing heat losses.

D. Magnitude of the synchronisation event

The magnitude of the synchronisation event, namely the duration of a power outage or the size of a collective step increment in the temperature set point of the devices, has a direct impact on the shape of the response. This magnitude determines the proportion of ACs that will change state (on to off or vice versa) with the event.

Let us compare the cases in Figure 7, showing the responses when all of the ACs are required to raise their temperature set-points by 25%, 50%, 75% and 100% of their hysteresis width. We see that larger steps (and, equivalently, longer power outages) produce larger oscillations as a higher proportion of devices is forced to change state. If the event is large enough to force all of the ACs to change to the same state, the normalised power response will reach its minimum (e.g., the 1 ^oC step response in Figure 7) or its maximum (e.g., a long enough power outage). Larger events do not affect further the amplitude of the saturated demand, but in general continue to have an impact on the time at which such demand will stay at this level (e.g., a step increment $u(t)$ of 3 ^oC will maintain all of the ACs off for longer than a 2 ^oC increment. Note that the magnitude of the event ceases to prolong this time when $\theta = \theta_m$ or $\theta = \theta_M$, where θ_m is the asymptotic temperature achieved by an AC if its compressor remains on indefinitely and θ_M is the asymptotic temperature when the compressor remains off (i.e., the outside temperature θ_a plus a correction factor due to internal heat loads). For instance, once the temperatures reach θ_M during a power outage, the duration of such outage is irrelevant to how long the devices will remain on when the power is restored.

The magnitude of the temperature set-point changes (but not outages) also affects the new steady state power demand.

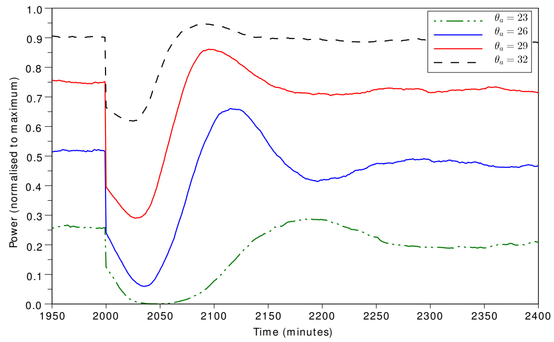

Fig. 8. Effect of the steady aggregate demand cycle on the aggregate response to the same external event (a 0.5 ^oC step). The different demands are obtained by considering different outside temperatures θ_a.

It can be seen in Figure 7 that the larger the increase in the reference temperature, the lower the steady state power demand. This is because the difference between the desired and outside temperature decreases as the ACs remain on for shorter periods of time. Conversely, lowering the temperature increases the ACs power use.

The magnitude of the event also affects the phase of the oscillations, as a larger event will cause that, on average, the ACs that were affected by the event take more time to change state. The frequency of the oscillations is independent of the magnitude of the event in the case of power outages, as when the electricity is restored, the rate at which the temperature changes within the hysteresis band is unrelated to the length of the outage. A step response, on the other hand, redefines the boundaries of the hysteresis band, which affects the rate at which the temperature varies within it.

E. Steady-state aggregate demand

Assuming steady state at the time of the synchronising event, different values of the aggregate demand at the time of the event will lead to different responses, even for the same event. For example, Figure 8 shows how different steady-state demands (obtained by varying the outside temperature θ_a) affect the aggregate response to the same event: a common 0.5 ^oC step change in the temperature set-points.

We observe that the oscillations are larger when the steady state demand is closer to 50%. Intuitively, this can be explained by the fact that there is "more room" for the demand to swing both upwards and downwards.

Note that for all of the scenarios in Figure 8 except for the highest aggregate demand ($\theta_a = 32$), a step change in the temperature set point equivalent to half the hysteresis width (0.5 ^oC) roughly halves the demand instantaneously. For high average duty cycles on the other hand (e.g., for $\theta_a = 32$), this is not the case, because a significant proportion of the ACs are *on and not within such band*. These are ACs that, despite operating at 100% duty cycle, are not capable of maintaining the temperature within the desired range (e.g., because they are undersized). These operating ACs are unaffected by the set point change, and continue in the on state after the step. This observation is of particular importance if DLC is implemented for load reduction on peak days (when ACs are likely to be operating at a high duty cycle) and the devices are being controlled through a temperature set point change.

V. Conclusions

When a population of ACs is subject to a synchronisation event such as a power outage, the aggregate power demand presents underdamped oscillations. This phenomenon has been reported in the literature for populations of ACs and other TCLs [2], [10], [14].

We have explained these oscillations in terms of the temperature distributions in the population and studied numerically how these oscillations depend on the characteristics of the population and the synchronisation event. The main findings of our study are (a) more heterogeneous populations present more decoherent responses (i.e., more damped oscillations); (b) heterogeneity does not affect the period of the oscillations; and (c) the type of distribution of the parameters R, P and C in the population does not have a significant impact in the characteristics of the oscillations.

We have successfully identified fundamental relations between the parameters of the ACs and their aggregate dynamics. Apart from aiding the understanding of fundamental behaviour, these relations can help develop and validate models that describe these dynamics [2], [3], [7].

References

[1] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, Jan 2011.

[2] D. 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.

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

[4] S. Bashash and H. Fathy, "Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads," in *Proc. American Control Conference*, 2011.

[5] 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 Conf.*, august 2011.

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

[7] C. Perfumo, E. Kofman, J. Braslavsky, and J. 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.

[8] N. A. Sinitsyn, S. Kundu, and S. Backhaus, "Safe protocols for generating power pulses with heterogeneous populations of thermostatically controlled loads," *Energy Conversion and Management*, vol. 67, 2013.

[9] C. Chong and A. Debs, "Statistical synthesis of power system functional load models," vol. 18, dec. 1979, pp. 264 –269.

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

[11] 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 International*, 2011.

[12] 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.

[13] N. Lu and S. Katipamula, "Control strategies of thermostatically controlled appliances in a competitive electricity market," in *IEEE Power Engineering Society General Meeting*, June 2005.

[14] R. Malhame and C. Chong, "Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system," *Automatic Control, IEEE Transactions on*, 1985.

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

[16] A. Glen, L. L., and D. J.H., "Computing the distribution of the product of two continuous random variables," *Computational Statistics & Data Analysis*, vol. 44, no. 3, pp. 451–464, 2004.