Perfumo, Christian; Braslavsky, Julio; Ward, John K.; Kofman, Ernesto.; "An analytical characterisation of cold-load pickup oscillations in thermostatically controlled loads". Published in Proceedings of the 3rd Australian Control Conference (AUCC) (Perth, W.A. 4-5 November, 2013) (2013)

Available from: http://dx.doi.org/10.1109/AUCC.2013.6697272

An analytical characterisation of cold-load pickup oscillations in thermostatically controlled loads

Cristian Perfumo[1], Julio Braslavsky[1,2], John K. Ward[1] and Ernesto Kofman[3]

***Abstract*— Large groups of thermostatically controlled loads can be controlled to achieve the necessary balance between generation and demand in power networks. When a significant portion of a population of thermostatically controlled loads is forced to change their on-off state simultaneously, the aggregate power demand of such population presents large, underdamped oscillations, a well-known phenomenon referred to by power utilities as "cold-load pickup". Characterising these oscillations and, in general, the aggregate dynamics of the population facilitates mathematical analysis and control design. In this paper we present a stochastic model for the power response and derive simple expressions for the period and envelope of the oscillations.**

I. Introduction

Thermostatically controlled loads (TCLs) such as air conditioners (ACs) have been identified as promising loads for demand response due to their widespread use and inherent energy storage capabilities [1]. By controlling large groups of these loads, the necessary balance between generation and consumption in electricity grids can be achieved from the demand side, potentially deferring expensive infrastructure augmentation and facilitate renewable energy integration.

However, remotely controlling groups of TCLs may result in large transients in the collective power demand, which might lead to undesired peaks in the demand. Figure 1, presents one of these situations, in which all of the devices are subject to a step change in their temperature set point at $t = 1000$. Similar oscillations are observed when, following a power outage, the power of all of the TCLs is restored, and when a load control (LC) event is finished and full independent control is returned to the TCLs while they are still substantially synchronised.

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 the aggregate demand. This phenomenon by which the synchronisation of TCLs produces large oscillations in their aggregate power demand is traditionally known as *cold-load pickup* [2]. As shown in Figure 1, this phenomenon presents damped oscillatory transients, which eventually converge to a particular value. We refer to this value as the *steady state aggregate demand* of the population. Regardless of the initial conditions, the heterogeneity in the population causes some of the TCLs to turn off (and then back on) sooner than others until the TCLs become *unsynchronised* (i.e., load diversity is restored) and the aggregate power demand settles at the steady state demand.

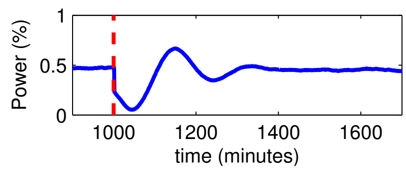

Fig. 1: Example synchronisation event that can cause oscillations in the aggregate power demand of a population of TCLs. At $t = 1000$ all of the TCLs are subject to a common step change in their temperature set points.

The exact shape of these oscillations depends on the characteristics of the TCLs, their operating conditions and the external event causing the synchronisation. This poses the question of how to characterise such response as a function of these parameters.

One alternative to achieve this characterisation is via System Identification [3]. This approach was applied successfully to populations of ACs in [1], using a series of temperature set point changes as the input signal, the aggregate power demand as the output signal and a first-order ARMAX model structure.

There are, however, two main drawbacks applying system identification to model populations of TCLs. Firstly, a new set of model parameters has to be numerically computed for every different population considered. Secondly, producing the input and output data required for the identification experiment may have undesirable consequences (e.g., it could disrupt user comfort beyond acceptable limits or cause undesirable peaks in the electricity network).

There are few analytic alternatives to system identification that do not result in overly complex models [1], [4], [5]. In particular, the work in [5] proposes a simple, second-order model for the aggregate power demand, assuming a 50% duty cycle in the population to achieve a robust nominal model.

In the present paper we generalise the approach in [5] to any duty cycle, and satisfactorily validate it against the oscillations observed by performing Monte Carlo simulations of the aggregate response for a population of 10,000 ACs. Specifically, we present a stochastic model for the power oscillations and derive simple expressions for their period and envelope. We show that the modelling approach from [5] readily generalises when the duty cycle assumption is

[1]CSIRO, Newcastle, NSW, Australia {cristian.perfumo, julio.braslavsky, john.k.ward}@csiro.au
[2]School of Elec. Eng. and Comp. Sci. University of Newcastle, Australia
[3]CIFASIS-CONICET, Department of Control, FCEIA-UNR Riobamba 245 bis (2000) Rosario, Argentina kofman@fceia.unr.edu.ar

relaxed, and allows the development of more accurate simple models for analysis, transient response prediction and tight control design when the population is not operating at its "central", 50% duty cycle point.

II. Assumed population and preliminary expressions

Assuming a population of n ACs, the dynamics of the i-th AC in the population can be modelled by the hybrid state model

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

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

presented in [6], where $\theta_i(t)$ is the room temperature, θ_a is the ambient temperature outside the rooms (°C), C_i and R_i are the i-th room thermal capacitance (kWh/°C) and thermal resistance (°C/kW), and P_i is the cooling thermal power of the i-th AC (kW). The binary variable $m_i \in \{0,1\}$ ($i \in \{1,2,\ldots,n\}$,) represents the state of the compressor which switches *on* the AC ($m_i = 1$) or *off* ($m_i = 0$) to maintain the temperature θ_i within the pre-specified hysteresis band $[\theta_i^-, \theta_i^+]$, centred at $\theta_i^{\mathrm{r}} = (\theta_i^- + \theta_i^+)/2$. The input signal w_i represents unpredictable thermal disturbances, and u is the proposed control signal (common to all ACs) to introduce small temporary temperature set-point offsets to the population during LC events.

Our objective in this paper is to obtain an analytical description of the response of the aggregate demand to a simultaneous step change in temperature set point in the whole population. To render this modelling problem tractable, let us start by defining a set of simplifying assumptions about the ensemble of ACs to be described:

- H.1 *Identical hysteresis bands:* All of the ACs in the population have the same set point temperature $\theta^{\mathrm{r}} = (\theta_+ + \theta_-)/2$ and the same hysteresis width $\theta_+ - \theta_- = 1$.
- H.2 *Uniform temperature distribution:* At $t = 0$, the temperatures are uniformly distributed in the interval $[\theta_-, \theta_+]$.
- H.3 *Triangular temperature evolution:* For each AC, the temperature changes at a *constant rate* when the AC is on and at a different *constant rate* when it is off.
- H.4 *Small disturbances:* The noise term $w(t)$ in (1) is negligible for each AC.
- H.5 *Distributed capacity only:* The parameter C is distributed in the population according to some probability distribution. The parameters P and R are identical for all the ACs.

Assumptions H.1–H.5 are required simplifications for our analysis. While they may seem overly restrictive, it has been shown that the controllers designed using these simplifying assumptions still preserve good performance even when these assumptions are relaxed [5].

The assumption of identical hysteresis bands (H.1) is considered purely for convenience, as it greatly simplifies the mathematical analysis. However, note that H.1 is without loss of generality, as the hysteresis band and the temperature set point of each device can be normalised.

Regarding the uniform temperature distribution assumption (H.2), the larger the average steady state duty cycle in the population, the less reasonable it is to assume a uniform temperature distribution because more and more ACs are unable to keep their temperatures within the hysteresis band even for a 100% duty cycle. Also, the steady state distribution of temperatures is affected by the level of heterogeneity: the less variability in the population, the more uniform the temperature distributions within the hysteresis band [1]. Thus, the uniform temperature distribution assumption becomes less reasonable for highly heterogeneous populations operating close to their maximum capacity.

The triangular temperature evolution assumption (H.3) implies that the rates of change in temperature both for $m(t) = 0$ and $m(t) = 1$ can be approximated as constants. This approximation has also been used in [7].

We neglect the disturbance $w(t)$ (H.4) 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. Additionally, more heterogeneity in the population (e.g., larger variance of $w(t)$) results in more damped oscillations, which means that a model developed assuming small $w(t)$ will constitute an "upper bound" of $D(t)$ for larger $w(t)$.

Finally, the assumption H.5, of capacity only being distributed, constitutes a trade-off between considering a heterogeneous population of ACs and a simplification that results in a tractable mathematical expression for our analysis.

Under assumptions H.1-H.5, we analyse how a population of ACs, each governed by (1)-(2), reacts when a step change of amplitude 0.5 (^{o}C) in the temperature set point is applied at $t = 0$, moving the boundaries of the hysteresis band as shown in Figure 2. We will refer to these new boundaries as $\theta_-^{post} = \theta_- + 0.5$ and $\theta_+^{post} = \theta_+ + 0.5$. We can see in Figure 2 that the ACs that were operating and had temperatures in $[\theta_-, \theta_-^{post}]$ before the step, switched off after it.

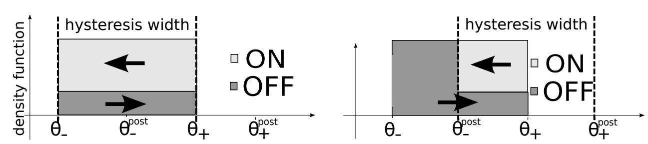

(a) Before step (b) After step

Fig. 2: Schematic distribution of temperatures before and after a 0.5^oC step in the set point (assuming H.1–H.5). The arrows indicate how the temperatures are moving.

Under assumptions H.1-H.5, the rate at which the temperature $\theta_i(t)$ increases for the i-th AC (as described in (1)) when the device is turned off is given by the constant v_i, defined by

$$v_i = \frac{\theta_a - \theta^{\mathrm{r}}}{C_i R}. \quad (3)$$

Let $T_{\mathrm{on}}^{\mathrm{i}}$ be the time it takes the i-th AC to bring its temperature $\theta_i(t)$ down from θ_+^{post} to θ_-^{post}. Analogously, let $T_{\text{off}}^{\text{i}}$ be the time it takes for the temperature to raise from θ_-^{post} to θ_+^{post} when the AC is turned off. Assuming a hysteresis width of 1 ^{o}C (H.1), we have that $T_{\text{off}}^{\text{i}} v_i = T_{\text{off}}^{\text{i}} \frac{\theta_a - \theta^{\text{r}}}{C_i R} = 1$ and, similarly, from (1) and under H.1-H.5 we have that $T_{\text{on}}^{\text{i}} \frac{\theta_a - \theta^{\text{r}} - RP}{C_i R} = -1$. Thus,

$$\frac{T_{\text{off}}^{\text{i}}}{T_{\text{on}}^{\text{i}}} = \frac{RP}{\theta_a - \theta^{\text{r}}} - 1, \tag{4}$$

which indicates that the ratio between $T_{\text{off}}^{\text{i}}$ and T_{on}^{i} is the same for all the ACs in the population (because of the assumption of distributed capacity only, H.5).

We then define the constants

$$x_{\text{off}} = T_{\text{off}}^{\text{i}} v_i = 1, \tag{5}$$

$$x_{\text{on}} = T_{\text{on}}^{\text{i}} v_i = \left(\frac{RP}{\theta_a - \theta^{\text{r}}} - 1\right)^{-1}, \tag{6}$$

$$x_{\text{c}} = x_{\text{on}} + x_{\text{off}} \tag{7}$$

and

$$D_{\text{c}} = \frac{x_{\text{on}}}{x_{\text{c}}} = \frac{T_{\text{on}}}{T_{\text{on}} + T_{\text{off}}}. \tag{8}$$

Note that in (8) we have dropped the indices in T_{on}^{i} and $T_{\text{off}}^{\text{i}}$ because D_{c} represents the average duty cycle in the population (in fact, because of H.1–H.5, D_{c} represents the duty cycle of *each* AC in the population).

Let us now introduce the variable

$$x_i(t) = x_i(0) + v_i t, \tag{9}$$

where

$$x_i(0) = \begin{cases} x_{\text{on}} + (\theta_i(0) - \theta_-^{post}) x_{\text{off}} & \text{if } m_i(0^-) = 0, \\ (\theta_+^{post} - \theta_i(0)) x_{\text{on}} & \text{if } m_i(0^-) = 1, \end{cases} \tag{10}$$

where $m_i(t)$ is the compressor state of the ith AC in the population, defined in (2).

Intuitively, we can say that $x_i(t)$ is the "unwrapped" mapping of the temperature $\theta_i(t)$ of the ith AC in the population. Figure 3 illustrates the equivalence between θ and x. We can see in Figure 3 that when the temperature $\theta_i(t) = \theta_+^{post}$ (the AC switches on), $x_i(t) = k x_{\text{c}}$ for some $k = 1, 2, \ldots$ Conversely, when $\theta_i(t) = \theta_-^{post}$ (the AC switches off), $x_i(t) = k x_{\text{c}} + x_{\text{on}}$. The same applies to $\theta_j(t)$ and $x_j(t)$. We will say that

- $x_i(t)$ is in an *on interval* whenever $x_i(t) \in [k x_{\text{c}}, k x_{\text{c}} + x_{\text{on}}]$, for $k = 1, 2, \ldots$, and that
- $x_i(t)$ is in an *off interval* whenever $x_i(t) \in [(k-1) x_{\text{c}} + x_{\text{on}}, k x_{\text{c}}]$ for $k = 1, 2, \ldots$.

Because of the assumption of distributed capacity only (H.5), note in Figure 3 that even though the duty cycles of the ith and jth ACs are the same in the population, the times they remain on and off are, in general, different (i.e., $T_{\text{on}}^{\text{i}}/(T_{\text{on}}^{\text{i}} + T_{\text{off}}^{\text{i}}) = T_{\text{on}}^{\text{j}}/(T_{\text{on}}^{\text{j}} + T_{\text{off}}^{\text{j}})$ even though $T_{\text{on}}^{\text{i}} \neq T_{\text{on}}^{\text{j}}$ and $T_{\text{off}}^{\text{i}} \neq T_{\text{off}}^{\text{j}}$).

Conversely, we can define a mapping from $x_i(t)$ to $\theta_i(t)$

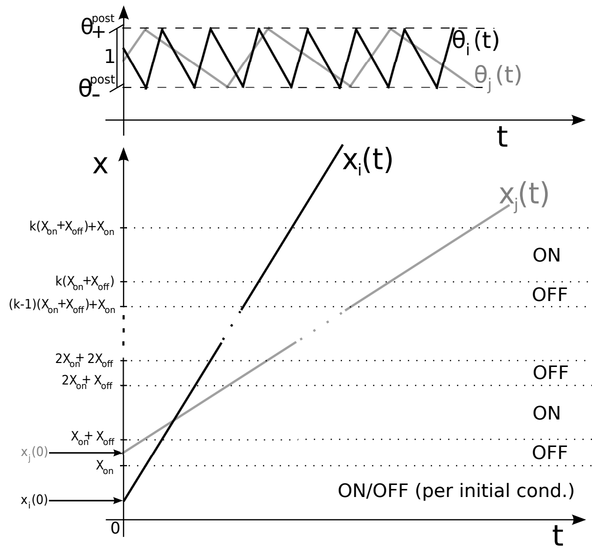

Fig. 3: Top: temperatures $\theta_i(t)$ and $\theta_j(t)$ of two ACs. Bottom: their "unwrapped" equivalents $x_i(t)$ and $x_j(t)$.

as

$$\theta_i(t) = \begin{cases} x_i(t) - x_{\text{c}}(x_i(t) \div x_{\text{c}}) + \theta_- & \text{if } m_i(0^-) = 0, \\ \theta_+ - (x_i(t) - x_{\text{c}}(x_i(t) \div x_{\text{c}}) - x_{\text{on}}) \frac{x_{\text{off}}}{x_{\text{on}}} & \text{if } m_i(0^-) = 1. \end{cases} \tag{11}$$

where $y \div z$ represents the integer part of the division y/z.

Lastly, for a large enough number n of ACs in the population, we can model the distribution of $\{x_1(t), x_2(t), \ldots, x_n(t)\}$ with a random variable $X(t)$.

Now we are ready to derive a stochastic expression that models the step response of a population of ACs to a common change in the temperature set point of the devices.

III. Stochastic characterisation of the aggregate power demand of a population of ACs

*Proposition 1 (***Characterisation of $D(t)$***): Let us consider a population of ACs where the dynamics of each device are described by* (1) *and* (2)*. Assuming H.1-H.5, if the temperature set points of all the devices in the population are raised by 0.5 ^{o}C at time $t = 0$, the probability $D(t)$ that a randomly-picked AC be operating at time t is given by:*

$$D(t) = \frac{\Pr[X(t) < x_{\text{on}}]}{2 + x_{\text{on}}/x_{\text{off}}} + \sum_{k=1}^{\infty} \Pr[X(t) < k x_{\text{c}} + x_{\text{on}}] - \Pr[X(t) < k x_{\text{c}}] \tag{12}$$

where x_{off}, x_{on} and x_{c} are defined as (5), (6) *and* (7) *respectively, $X(t)$ is a random variable modelling the distribution of $\{x_1(t), x_2(t), \ldots, x_n(t)\}$, $x_i(t)$ is defined in* (9)*, and $\Pr[\cdot]$ is the probability operator.*

The proof of Proposition 1 can be found in Section A of the Appendix.

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 *on* interval (i.e., $x_i(t) \in [k x_{\text{c}}, k x_{\text{c}} + x_{\text{on}}]$, for $k \geq 1$) plus a term divided by a correction factor $2 + x_{\text{on}}/x_{\text{off}}$

to take into account an anomaly in the first on interval (for details, see proof in Section A of the Appendix).

Note that for a large enough population, the probability (12) 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.5), this proportion is equal to the power consumption of all of the devices operating at time t, normalised by the maximum demand. We deliberately use the notation $D(t)$ to refer to the probability of a randomly-selected AC be operating, the proportion of operating ACs in the population and the normalised power consumption, unless stated otherwise.

In order to calculate actual values for (12), it is necessary to know how the parameter C is distributed in the population. A corollary of Proposition 1 follows from considering C log-normally distributed in the population of TCLs. We adopt a log-normal distribution (as done in [1]) because it is suitable for non-negative parameters and has a complexity of description that is only moderate.

Corollary 1: **Characterisation of $D(t)$ for lognormally distributed parameters**: *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 (Equation* (3)*) is log-normally distributed in the population with standard deviation σ_v and mean μ_v satisfying the ratio*

$$\sigma_r = \frac{\sigma_v}{\mu_v} = \frac{\sigma_C}{\mu_C}. \tag{13}$$

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

$$\begin{aligned} D(t) \approx{}& \frac{\alpha}{2} + \frac{\alpha}{2}\mathrm{erf}\left[\frac{\log(x_{\mathrm{on}}) - \tau}{\sqrt{2}\sigma_r}\right] \\ &+ \frac{1}{2}\sum_{k=1}^{\infty}\mathrm{erf}\left[\frac{\log(kx_{\mathrm{c}} + x_{\mathrm{on}}) - \tau)}{\sqrt{2}\sigma_r}\right] - \mathrm{erf}\left[\frac{\log(kx_{\mathrm{c}}) - \tau}{\sqrt{2}\sigma_r}\right], \end{aligned} \tag{14}$$

where

$$\mu_x(0) = \frac{7}{8}x_{\mathrm{on}}(x_{\mathrm{on}}/x_{\mathrm{off}} + 1)(1 - D_{\mathrm{c}}), \tag{15}$$
$$\tau = \log(\mu_x(0) + \mu_v t),$$

$\mu_x(t)$ is the mean of the values $x_i(t)$ at time t, D_{c} is defined in (8)*, $\alpha = \frac{1}{2 + x_{\mathrm{off}}/x_{\mathrm{on}}}$ and $\mathrm{erf}[\cdot]$ is the Gauss error function.*

See the proof of Corollary 1 in the Appendix of [8].

Figures 4, 5 and 6 plot in dashed lines the expected power response according to (14) to a 0.5 ^oC step at $t = 0$. The output is normalised to the maximum power output (all ACs turned on). The figures also show, as a solid line, the output to the same input when we simulate the response of 10000 ACs (using (1)-(2)) assuming C distributed log-normally in the population and P and R not distributed. The only difference between Figures 4, 5 and 6 is the standard deviation to mean ratio σ_{r} (0.05, 0.1 and 0.2 respectively) of the parameter C.

The parameters used for our simulations are based on those presented in [1], most of which are derived from a list of references in the literature that investigate thermal properties of buildings. One exception is the hysteresis width, which was considered 0.5 ^oC in [1] but we double to 1 ^oC here, as it has more typically been observed in practice. Table I summarises our simulation parameters.

The ambient temperature was adjusted in each case to obtain a different duty cycle. For a given duty cycle D_c and a given reference temperature set point θ^{r}, the ambient temperature can be computed as

$$\theta_a(D_c) = \theta^{\mathrm{r}} + D_c RP. \tag{16}$$

Parameter	Value	Description
R	2 ^oC/kW	Thermal resistance
C	10 kWh/^oC	Mean thermal capacitance
P	14 kW	Thermal power
θ_-	19.5 ^oC	Lower end of hysteresis band
θ_+	20.5 ^oC	Higher end of hysteresis band
σ_w	0.01	Standard deviation of the noise process w in Eq. (1)
σ_{r}	0.05/0.1/ 0.2/0.5	Ratio between standard deviation and mean thermal capacity C

TABLE I: Simulation parameters.

We observe in Figures 4 and 6 that the expression (14) agrees with the simulated $D(t)$ most closely for medium duty cycles (e.g., $D_{\mathrm{c}} = 0.5$). Additionally, very high duty cycles (e.g. 0.9) present better fit that very low ones (e.g. 0.2) in terms of amplitude of the oscillations. This is due to the nature of the fist term in (12), which becomes a worst estimation of the initial state of the on and off populations when the duty cycle is lower than 0.5 (see Figure 9 in the Appendix). The frequency of the simulated $D(t)$ is captured by (14) in all cases.

The following proposition uses (14) to analytically characterise the oscillations in the aggregate power demand caused by a step change in the temperature set point of the devices.

Proposition 2: **Period and envelope of oscillations in $D(t)$** *Under the assumptions of Corollary 1, the initial transients in the response of the aggregated power $D(t)$ to a step change in temperature set point display damped oscillations with period*

$$T \approx \frac{x_{\mathrm{on}} + x_{\mathrm{off}}}{\mu_v}, \tag{17}$$

within an envelope that can be approximated by

$$1 - \mathrm{erf}(x_{\mathrm{off}}/z(t)) \leq D(t) \leq \mathrm{erf}(x_{\mathrm{on}}/z(t)), \tag{18}$$

where

$$z(t) = 2\sqrt{2}\sigma_r(\mu_x(0) + \mu_v t - x_{\mathrm{on}}/2), \tag{19}$$

$$\mu_v = \frac{\theta_a - \theta^r}{R}e^{(-\tilde{\mu_C} + \tilde{\sigma_C}^2/2)}, \tag{20}$$

considering that $log(C) \sim \mathcal{N}(\tilde{\mu_C}, \tilde{\sigma_C}^2)$. The definitions of x_{on}, x_{off} and $\mu_x(0)$ are given in (6)*,* (5) *and* (15) *respectively.*

For the proof of Prop. 2 see the the Appendix in [8].

Note in (17) and (18) that because $\mathrm{erf}(x_{\mathrm{on}}/z(t))$ and $\mathrm{erf}(x_{\mathrm{off}}/z(t))$ decrease monotonically from 1 to 0 for pos-

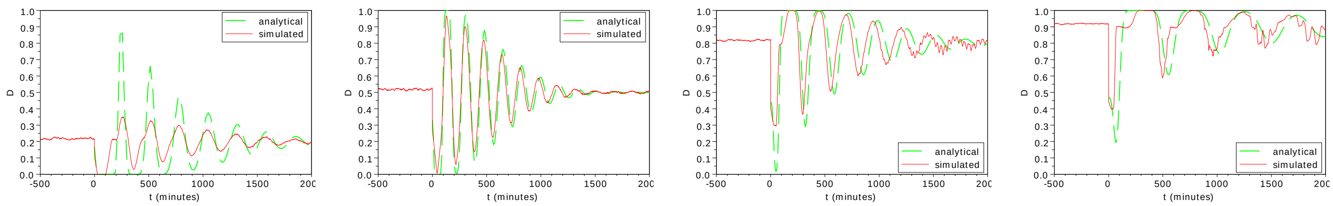

(a) Duty cycle = 0.2 (b) Duty cycle = 0.5 (c) Duty cycle = 0.8 (d) Duty cycle = 0.9

Fig. 4: For $\sigma_{\mathrm{r}} = 0.05$, analytical (as per (14)) and simulated D (normalised power demand) response to a common 0.5 ^oC step change in temperature set point. Simulated results: 10000 ACs with log-normally distributed C, and constant R and P for all the ACs, with parameters from Table I. The different duty cycles were obtained by setting θ_a according to (16).

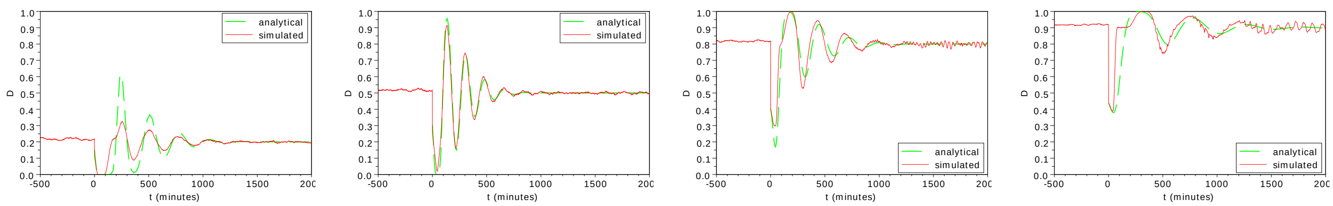

(a) Duty cycle = 0.2 (b) Duty cycle = 0.5 (c) Duty cycle = 0.8 (d) Duty cycle = 0.9

Fig. 5: Same as Figure 4, except for $\sigma_{\mathrm{r}} = 0.1$.

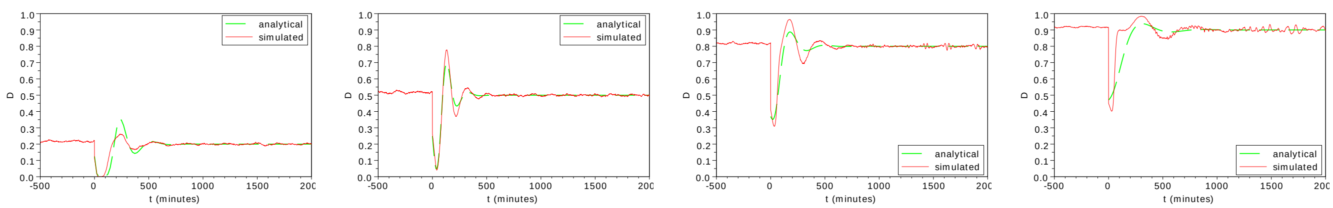

(a) Duty cycle = 0.2 (b) Duty cycle = 0.5 (c) Duty cycle = 0.8 (d) Duty cycle = 0.9

Fig. 6: Same as Figure 4, except for $\sigma_{\mathrm{r}} = 0.2$.

itive values of z, the amplitude in the oscillations in $D(t)$ decreases over time. Moreover, as we can see in (18) and (19), the larger σ_{r} (more heterogeneity in the population), the faster the envelope decreases, as observed in [5].

Figure 7 depicts the envelope bound (18) of the oscillations in the step power response along with the simulated power output of 10000 ACs (Table I) for different values of σ_{r} as a function of the offset and rescaled time z, as defined in (19). Plotting as a function of z allows us to show responses for different values of σ_{r} within a common envelope. Each subplot represents a different duty cycle. The curves start at the corresponding value of z that makes $t = 0$. We can see in the curves that (18) is a poorer approximation of the envelope for duty cycles far from 0.5 (an especially for low duty cycles, where the computed envelope is too conservative).

Figure 8 depicts the same envelope (18), but the power responses within it are the ones approximated with (14). By comparing Figures 7 and 8, we can see why (18) is too conservative for low duty cycles: as shown in Figures 4, 5 and 6, (14) approximates $D(t)$ poorly for these average duty cycles.

IV. Conclusions

The step response of the aggregate power of a population of ACs under the assumptions H.1-H.5 is dominated by decaying oscillations, which corroborates the simulation results reported by a number of authors [2], [9], [10], [1]. We have characterised these oscillations 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 evolution of this proportion of operating ACs is directly related to the their aggregate demand (and directly propor-

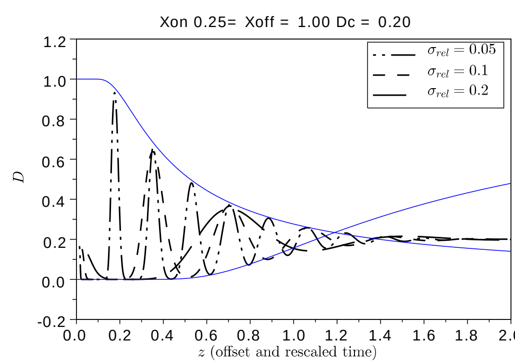

Fig. 8: Envelope of power peaks for 20% duty cycle. Dotted lines: normalised power demand D for different values of σ_{r} according to (14), adjusted to the offset and rescaled time variable z. Continuous line: envelope as per (18).

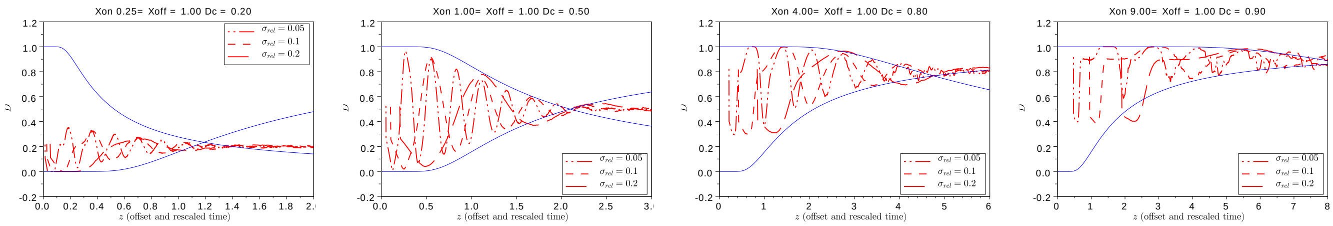

(a) Duty cycle = 0.2 (b) Duty cycle = 0.5 (c) Duty cycle = 0.8 (d) Duty cycle = 0.9

Fig. 7: Envelope of power peaks for different duty cycles. Dotted lines: normalised power demand D of 10000 ACs simulated according to (1)-(2) for different values of $\sigma_{\rm r}$, adjusted to the time variable z. Continuous line: envelope given by (18).

tional if all of the ACs have the same electrical power), and therefore is also captured in $D(t)$. Our model satisfactorily captures the dynamics of a numerically simulated population of ACs for medium and large steady state demands.

In spite of its complexity, we successfully used this stochastic model to derive simple mathematical approximations of the period and amplitude envelope the response. These expressions are of particular interest when assessing the implications of using TCLs for demand response.

Appendix

A. Proof of Proposition 1 (Characterisation of $D(t)$)

Proof: From (9) and (10), the operational state $m_i(t)$ of the i-th AC in the population is defined by

$$m_i(t) = \begin{cases} 0 & \text{if } (k-1)x_{\rm c} + x_{\rm on} \le x_i(t) < kx_{\rm c} \\ 1 & \text{if } kx_{\rm c} \le x_i(t) < kx_{\rm c} + x_{\rm on}, \end{cases} \tag{21}$$

for $k = 1, 2, \ldots$ Thus, for a number of ACs sufficiently large, the probability that a randomly picked AC be operating at time t given $X(t) > x_{\rm on}$ is

$$\Pr[m(t) = 1 | X(t) > x_{\rm on}] = \sum_{k=1}^{\infty} \Pr[X(t) < kx_{\rm c} + x_{\rm on}] - \Pr[X(t) < kx_{\rm c}]. \tag{22}$$

The fraction of ACs that are on when $x_i(t) < x_{\rm on}$ is calculated in a different way. When (10) is applied to every point in the temperature distributions before and after the step change shown in Figure 2(b), the distribution of the initial values of $X(t)$ for all of the ACs, namely $X(0^-)$ and $X(0^+)$, is as shown in Figure 9.

Note that only a fraction α of the ACs that satisfy $x_i(0) < x_{\rm on}$ is turned on. Let A_j be the area of the rectangle number j in Figure 9. Then

$$\alpha = \frac{A_1}{A_1 + A_2 + A_4} = \frac{A_1}{2A_1 + A_2} = \frac{A_1}{2A_1 + A_1 x_{\rm off}/x_{\rm on}} = \frac{1}{2 + x_{\rm off}/x_{\rm on}}. \tag{23}$$

Assuming H.1-H.5, the probability of a randomly-picked AC operating at time t given $X(t) < x_{\rm on}$ is

$$\Pr[m(t) = 1 | X(t) < x_{\rm on}] = \alpha \Pr[X(t) < x_{\rm on}]. \tag{24}$$

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,

$$D(t) = \alpha \Pr[X(t) < x_{\rm on}] + \sum_{k=1}^{\infty} \Pr[X(t) < kx_{\rm c} + x_{\rm on}] - \Pr[X(t) < kx_{\rm c}].$$

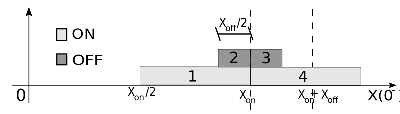

(a) Before step

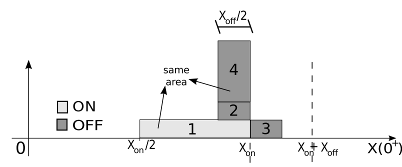

(b) After step

Fig. 9: Distribution of $X(0^-)$ and $X(0^+)$ (before and after the temperature set point step change at $t = 0$). 1: ON before and after the step; 2 and 3: OFF before and after the step; 4: ON before the step, OFF after it (same area as 4).

■

References

[1] 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.

[2] 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.

[3] L. Ljung, *System identification: theory for the user*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1986.

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

[5] 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.

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

[7] 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.

[8] 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.

[9] C. Chong and R. Malhame, “Statistical synthesis of physically based load models with applications to cold load pickup,” *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-103, no. 7, pp. 1621 –1628, jul. 1984.

[10] R. Malhame, “A statistical approach for modeling a class of power system loads,” Ph.D. dissertation, Georgia Institute of Technology, 1983.