

Low-complexity control algorithm for decentralised demand response using thermostatic loads

Tindemans, Simon; Strbac, Goran

DOI

[10.1109/EEEIC.2019.8783359](https://doi.org/10.1109/EEEIC.2019.8783359)

Publication date

2019

Document Version

Final published version

Published in

Proceedings - 2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe, EEEIC/I and CPS Europe 2019

Citation (APA)

Tindemans, S., & Strbac, G. (2019). Low-complexity control algorithm for decentralised demand response using thermostatic loads. In *Proceedings - 2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe, EEEIC/I and CPS Europe 2019* (pp. 1-6). Article 8783359 IEEE. <https://doi.org/10.1109/EEEIC.2019.8783359>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Low-complexity control algorithm for decentralised demand response using thermostatic loads

Simon Tindemans

Department of Electrical Sustainable Energy
Delft University of Technology
Delft, The Netherlands
s.h.tindemans@tudelft.nl

Goran Strbac

Department of Electrical and Electronic Engineering
Imperial College London
London, United Kingdom
g.strbac@imperial.ac.uk

Abstract—Thermostatically controlled loads such as refrigerators are exceptionally suitable as a flexible demand resource. This paper derives a decentralised load control algorithm for refrigerators. It is adapted from an existing continuous time control approach, with the aim to achieve low computational complexity and an ability to handle discrete time steps of variable length – desirable features for embedding in appliances and high-throughput simulations. Simulation results of large populations of heterogeneous appliances illustrate the accurate aggregate control of power consumption and high computational efficiency.

Keywords—thermostatically controlled loads, demand response, decentralized control, aggregate control

I. INTRODUCTION

The physical characteristics of refrigerators and other thermostatically controlled loads (TCLs) make them exceptionally suitable as a low-cost provider of flexibility to the grid: their power consumption can be shifted by tens of minutes without noticeable effects on cooling performance. This flexibility can then be used for the provision of response and reserve services, to reduce extreme load levels and to alleviate ramping constraints [1]. Given the large number of devices involved, large populations can be effectively controlled in a decentralised manner using randomised control schemes, as proposed e.g. in [2]–[4]. Practical implementation must also consider constraints on implementation, computation and communication, as discussed e.g. in [5], [6].

A robust decentralised control scheme for heterogeneous TCLs was introduced in [7], and extended in [8] to allow for scenarios in which TCLs collectively absorb energy from the grid. The control strategy has the desirable feature that it requires only one-way broadcast information, yet achieves tracking of a reference signal that is exact in expectation (i.e. exact for large numbers) without violating temperature limits in individual devices. However, its continuous time formulation in integral form was not conducive to implementation on embedded controllers or for rapid simulation of many devices at once.

This paper addresses that shortcoming by deriving a discrete-time control algorithm that implements the control strategy that was presented in [7], [8]. The algorithm is particularly

This research was supported by EDF Energy R&D UK Centre as part of the STAMINA project.

suitable for implementation on devices with computational constraints. Specifically, it avoids numerical integration and uses greedy time steps of variable size, so that real-time performance requirements are relaxed. First, a discretisation procedure is described, for the case of a piecewise-constant control signal, and the controller is reformulated in natural coordinates. Then, expressions are derived for each of the on/off-switching processes governing the behaviour of the device. Finally, an explicit control algorithm is provided, and its efficacy is illustrated with Python-based simulations for a heterogeneous population of refrigerators.

II. PRELIMINARIES

A. Appliance model

Throughout this paper, we consider the first order TCL model, expressed by the following differential equation for the temperature T^a of the compartment of appliance a :

$$\frac{dT^a(t)}{dt} = -\alpha^a [T^a(t) - T_{\text{off}}^a + c^a(t) \cdot (T_{\text{off}}^a - T_{\text{on}}^a)] \quad (1)$$

For ease of exposition, we shall refer to refrigerating appliances throughout, although the same model and control strategy can be used for other TCLs, e.g. space heaters. Here, $c^a(t) \in \{0, 1\}$ is the state of the compressor, and T_{off}^a and T_{on}^a are the asymptotic temperatures in the *off* and *on* states, respectively. The power consumption $P^a(t)$ of the appliance is assumed to be dominated by the compressor power consumption P_{on}^a , so that

$$P^a(t) = P_{\text{on}}^a c^a(t). \quad (2)$$

In the steady state (no control actions), the appliance is subject to a hysteresis controller that switches to the *on* state whenever an upper temperature bound T_{max}^a is reached, and to the *off* state when T_{min}^a is reached. This results in a periodic cycling of the power consumption, with an average power level P_0^a . Let us consider that each appliance a has a model $\mathcal{M}^a = \{\alpha^a, P_{\text{on}}^a, T_{\text{off}}^a, T_{\text{on}}^a, T_{\text{min}}^a, T_{\text{max}}^a\}$ that is known to us, but an internal state $\mathcal{S}^a = \{T^a(t), c^a(t)\}$ that is not. We assume the latter is distributed according to a steady state distribution. Then, the power consumption is in effect a random process, with the expectation (at each time t)

$$E_{\mathcal{S}^a}[P^a(t)] = P_0^a. \quad (3)$$

B. Aggregate power modulation

The objective of the TCL demand response controller is to control the aggregate power consumption

$$P(t) = \sum_{a \in \mathcal{A}} P^a(t) \quad (4)$$

of a collection of appliances \mathcal{A} . In this paper we consider the control approach introduced in [7], which modulates the power consumption using a broadcast reference signal $\Pi(t)$ in the following way. Each device individually adapts its power consumption to $\Pi(t)$ in order to satisfy

$$E_{\mathcal{S}^a}[P(t)] = \Pi(t)P_0^a. \quad (5)$$

Clearly, $\Pi(t) = 1$ represents the steady state (3) and changes in $\Pi(t)$ are immediately reflected in the expected power consumption. Moreover, the controller maintains independence between appliances (conditional on the control signal $\Pi(\cdot)$) so that the central limit theorem can be applied to the total power consumption (4), resulting in

$$P(t) = \Pi(t) \sum_{a \in \mathcal{A}} P_0^a + O(\sqrt{|\mathcal{A}|}), \quad (6)$$

where the last term is a random process that decreases in relative importance to the first term as the set of appliances increases. We note that this is the case even for heterogeneous appliances.

The ability to closely track a reference signal was first demonstrated in [7]. In [8], the control signal was generated using a mixture of off-line scheduling and real-time control, and in [9] various frequency-sensitive controllers to locally compute $\Pi(t)$ (e.g. a simple droop controller) were implemented.

C. Distribution-referred control

The control approach introduced in [7] and extended in [8] is *distribution-referred* approach, because it implements control via the probability distribution of temperatures for appliances with a known model \mathcal{M}^a and unknown state \mathcal{S}^a . It defines a family of alternative distributions $f_z(T; \mathcal{M}^a)$ that varies continuously in the parameter z , containing as a special case the steady state temperature distribution [7, Eqs. (31)-(32)]

$$f_0(T; \mathcal{M}^a) = \frac{k^a}{(T_{\text{off}}^a - T)(T - T_{\text{on}}^a)} \quad (7)$$

with

$$k^a = \frac{T_{\text{off}}^a - T_{\text{on}}^a}{\log \left(\frac{(T_{\text{max}}^a - T_{\text{on}}^a)(T_{\text{min}}^a - T_{\text{off}}^a)}{(T_{\text{min}}^a - T_{\text{on}}^a)(T_{\text{max}}^a - T_{\text{off}}^a)} \right)} \quad (8)$$

The controller consists of two major elements, which are evaluated by each appliance in order:

- 1) Choose the distribution parameter $z(t)$ such that the power consumption tracks the reference signal $\Pi(t)$ according to (5). Determine the collective device switching actions required to keep device temperatures aligned with $f_{z(t)}(T)$, and identify the temperature limits $T_{\text{min}}(t) \geq T_{\text{min}}$ and $T_{\text{max}}(t) \leq T_{\text{max}}$.

- 2) Based on the *actual* appliance state \mathcal{S}^a , compute stochastic control actions, in the form of *on/off* switching. Switching events can be initiated in three distinct ways:

- Deterministic switching when temperature limits $T_{\text{min}}(t)$ or $T_{\text{max}}(t)$ are exceeded.
- Continuous-time stochastic switching at intermediate temperatures in order to shape the temperature distribution.
- Instantaneous stochastic switching on discrete changes of power setpoints, or when the controller switches between energy-provision and energy-absorption modes [8].

The temperature distribution and appliance switching phases for the discrete time control strategy are addressed in Sections III and IV, respectively. We will henceforth drop the appliance superscript a , because the control steps are executed locally within each appliance (or independently for each appliance in a simulation). Note that this implies a single model \mathcal{M} is used in the derivations, but the results remain valid for portfolios of heterogeneous devices (each with their own model).

D. Discretisation procedure

Moving from a continuous time formulation to a discrete time formulation, we consider a partitioning of the timeline by the ordered sequence of times $\{t_i\}$, indexed by the integer i , at which the controller is invoked. They define time intervals $(t_{i+1}, t_i]$ with durations $\Delta t_i = t_i - t_{i-1}$. Note that the duration Δt_i refers to the interval *prior* to t_i , and the intervals may have variable size. The reference signal $\Pi(t)$ is assumed to be piecewise constant, defined by

$$\Pi(t) = \Pi_i, \quad \text{for } t \in (t_{i-1}, t_i]. \quad (9)$$

The controller thus receives at t_i a new reference power level Π_{i+1} that must be applied for the upcoming interval $(t_i, t_{i+1}]$.

Although the discontinuous changes of reference power will trigger switching events at t_i , the other switching events may occur at any time t . In the discretised approximation of the continuous time controller, they will be synchronised with the control execution times t_i as follows. It is assumed that switching is immediate (at t_i).

- A violation of the temperature limits will trigger corrective switching as soon as it is detected.
- Switching due to toggling between energy-provision and energy-absorption modes is implemented as soon as a change in regime is detected.
- Continuous time stochastic switching is implemented by approximating the integrated switching rate (i.e. the switching probability) over Δt_i using the trapezoidal method, and executing any switching events at t_i (the end of the interval).

The algorithm is thus implemented in a ‘backward’ fashion, meaning that at time t_i , the algorithm implements switching actions resulting from reference changes at t_i , and those accumulated over the *preceding* interval $(t_{i-1}, t_i]$. The advantage of this approach is that the interval Δt_i can be chosen

opportunistically: the controller does not need to know in advance when the next time step will take place. This is convenient, for example when computational limitations cause a delay in intended invocation time, or when the time step adapts to sudden changes in grid frequency. It should be pointed out that this ‘backward’ integration does not delay the response to changes in reference power, which is implemented immediately at t_i .

III. DISTRIBUTION BEHAVIOUR

This section focuses on the first part of the controller. It computes the desired evolution of the probability distribution of temperatures of fridges with model \mathcal{M} , when tracking a piecewise constant reference Π_i . The derivation is initially performed in continuous time. The results are subsequently expressed in natural coordinates and restated in a form that is suitable for discrete-time evaluation.

A. Aggregate physics

The average temperature of a TCL population is affected by the desired power consumption $\Pi(t)$ according to [7, Eq. (26)]. With the convention that $t_{-1} = -\infty$ and $\Pi_0 = 1$ (assuming an initial steady state), it follows that

$$\bar{T}(t_i) = T_{\text{off}} - \alpha(T_{\text{off}} - \bar{T}_0) \sum_{j=0}^i \Pi_j \int_{t_{j-1}}^{t_j} e^{-\alpha(t_i-t')} dt', \quad (10)$$

where the steady state average temperature \bar{T}_0 is computed using [7, Eqs. (23) and (32)-(33)] as

$$\bar{T}_0 = T_{\text{off}} - k \times \log \left(\frac{T_{\text{max}} - T_{\text{on}}}{T_{\text{min}} - T_{\text{on}}} \right) \quad (11)$$

with k defined in (8). We define the dimensionless variable

$$z(t) = \frac{\bar{T}_0 - \bar{T}(t)}{T_{\text{off}} - \bar{T}_0}, \quad (12)$$

to parameterise the distributions $f_{z(t)}(T)$, and simplify the notation in what follows. Note that it is related to the σ variable used in [8] as $z = \sigma - 1$.

B. Controller modes

The algorithm in [7] implicitly generates the family of temperature distributions $f_{z(t)}(T)$ by the net heating rate $v(T, t)$, which is determined by averaging over devices in the `off` (heating) and `on` (cooling) states at time t and temperature T . The heating rate is controlled by a parameter $\beta(t)$ through $v(T, t) = \alpha\beta(t)(T - T_{\text{max}})$. The effect of this heating rate profile is a temperature distribution that contracts to the pivot temperature T_{max} in order to provide energy to the grid - and reverses this process to recover the energy supplied. In [8] it was coupled to a ‘mirrored’ controller that is capable of absorbing energy from the grid by contracting to the pivot temperature T_{min} . Switching between the two controller modes takes place whenever $\bar{T}(t)$ crosses \bar{T}_0 (when $z(t)$ crosses 1).

A generalised formulation covering both regimes is obtained by defining a heating rate of the form $v(T, t; R) =$

$\alpha\beta(t; R)(T - R(t))$, where $R \in \{T_{\text{min}}, T_{\text{max}}\}$ is a reference temperature, which acts as a pivot temperature for the controller, with the property $v(R, t; R) = 0$. The reference temperature is defined as follows:

$$R(t) = \begin{cases} T_{\text{max}}, & \text{if } \bar{T}(t) \geq \bar{T}_0 \\ T_{\text{min}}, & \text{if } \bar{T}(t) < \bar{T}_0 \end{cases} \quad (13)$$

C. Control parameter

The control parameter $\beta(t; R)$ is determined by the desired reference power $\Pi(t)$ according to [7, Eq. (36)]:

$$\begin{aligned} \beta(t; R) &= \frac{\Pi(t)(T_{\text{off}} - \bar{T}_0) - (T_{\text{off}} - \bar{T}(t))}{R(t) - \bar{T}(t)} \\ &= \frac{(\Pi(t) - 1) - z(t)}{z(t) - \zeta(R(t))} \end{aligned} \quad (14)$$

where

$$\zeta(R) = \frac{\bar{T}_0 - R}{T_{\text{off}} - \bar{T}_0}. \quad (15)$$

The denominator in the definition of β reflects, in dimensionless form, the energy limits of the TCL aggregate. Note also that β switches sign depending on the value of $R(t)$.

D. Distribution scaling

The controller has the effect of scaling the steady state temperature distribution $f_0(T)$ around the pivot temperature $R(t)$, such that the distribution does not exceed the temperature bounds T_{min} and T_{max} [7]. The extent of this scaling at time t_i is compactly represented by the scale parameter

$$\begin{aligned} s(t) &= \frac{R(t) - \bar{T}(t)}{R(t) - \bar{T}_0} \\ &= 1 - z(t)/\zeta(R(t)). \end{aligned} \quad (16)$$

E. Discretisation

We now consider the restriction of the continuous time controller to the set of discrete times t_i . We replace the coordinate $z(t)$ by its discretisation $z_i = z(t_i)$, which is computed from (10) as

$$z_i = \sum_{j=0}^i (\Pi_j - 1) \left(e^{-\alpha(t_i-t_j)} - e^{-\alpha(t_i-t_{j-1})} \right).$$

Updates to z_i are efficiently implemented using $z_0 = 0$ (for a steady state initialisation) and the recursive relation

$$z_i = z_{i-1} e^{-\alpha\Delta t_i} + (\Pi_i - 1)(1 - e^{-\alpha\Delta t_i}). \quad (17)$$

The discretised controller switches modes only at instants t_i , so $R(t)$ is approximated by the delayed function

$$\hat{R}(t) = R_i, \quad \text{for } t \in (t_{i-1}, t_i].$$

with

$$R_{i+1} = \begin{cases} T_{\text{max}}, & \text{if } z_i \leq 0, \\ T_{\text{min}}, & \text{if } z_i > 0. \end{cases} \quad (18)$$

Because our analysis focuses on the control time t_i , where $\hat{R}(t)$ and $\Pi(t)$ are potentially discontinuous, we introduce \pm -notation for the left and right limits at t_i :

$$R_i^- = \lim_{\varepsilon \downarrow 0} \hat{R}(t - \varepsilon) = R_i, \quad (19a)$$

$$R_i^+ = \lim_{\varepsilon \downarrow 0} \hat{R}(t + \varepsilon) = R_{i+1}. \quad (19b)$$

Similar definitions using left and right limits naturally apply to $\zeta(R)$, $s(t)$ and $\beta(t; R)$:

$$\zeta_i^\pm = \frac{\bar{T}_0 - R_i^\pm}{T_{\text{off}} - \bar{T}_0}, \quad \beta_i^- = \frac{(\Pi_i - 1) - z_i}{z_i - \zeta_i^-}, \quad (20a)$$

$$s_i^\pm = 1 - z_i/\zeta_i^\pm. \quad \beta_i^+ = \frac{(\Pi_{i+1} - 1) - z_i}{z_i - \zeta_i^+}, \quad (20b)$$

F. Energy and power constraints

The ability of the aggregate appliances to sustain a low or high power level is determined by operating temperature bounds of the appliance, applied to the distribution-averaged temperature: $\bar{T}(t) \in (T_{\min}, T_{\max})$ (no feasible solutions for the distribution $f_z(T)$ exists outside of this domain). However, because operation near the limits is infeasible in practice, due to diverging switching rates, we shall use a restricted range of operating temperatures that is scaled with a fraction $w < 1$ around the steady state operating temperature \bar{T}_0 :

$$(1 - w)\bar{T}_0 + wT_{\min} \leq \bar{T}(t) \leq (1 - w)\bar{T}_0 + wT_{\max}.$$

Rewriting this in terms of $z(t)$ and $\zeta(\cdot)$, we get

$$w\zeta(T_{\max}) \leq z(t) \leq w\zeta(T_{\min}).$$

Small excursions out of this temperature band will be permitted, but if this happens, the requested power level Π_{i+1} will be restricted to not exacerbate the excursion, using the relation (17). This leads to the update rule for Π_{i+1} :

$$\Pi_{i+1} := \begin{cases} \max(\Pi_{i+1}, 1 + w\zeta(T_{\max})), & \text{if } z_i \leq w\zeta(T_{\max}) \\ \min(\Pi_{i+1}, 1 + w\zeta(T_{\min})), & \text{if } z_i \geq w\zeta(T_{\min}) \\ \Pi_{i+1}, & \text{otherwise} \end{cases} \quad (21)$$

In addition to energy constraints related to the distribution-averaged temperature, the controller is subject to instantaneous power constraints that result from the maximum rate of change of the distribution. If the controller is in *energy provision mode* ($z_i \leq 0$), these power constraints are given by [8]

$$\left(\frac{\bar{T}_0 - T_{\min}}{T_{\max} - T_{\min}} \right) \left(\frac{T_{\text{off}} - T_{\max}}{T_{\text{off}} - \bar{T}_0} \right) \leq \Pi_{i+1} \leq \left(\frac{T_{\text{off}} - T_{\max}}{T_{\text{off}} - \bar{T}_0} \right) + \frac{(T_{\max} - \bar{T}_0)(T_{\max} - T_{\text{on}})}{(T_{\max} - T_{\min})(T_{\text{off}} - \bar{T}_0)}. \quad (22)$$

If the controller is in *energy absorption mode* ($z_i > 0$), the power constraints are given by

$$\left(\frac{T_{\max} - \bar{T}_0}{T_{\max} - T_{\min}} \right) \left(\frac{T_{\text{off}} - T_{\min}}{T_{\text{off}} - \bar{T}_0} \right) \leq \Pi_{i+1} \leq \left(\frac{T_{\text{off}} - T_{\min}}{T_{\text{off}} - \bar{T}_0} \right) + \frac{(\bar{T}_0 - T_{\min})(T_{\min} - T_{\text{on}})}{(T_{\max} - T_{\min})(T_{\text{off}} - \bar{T}_0)}. \quad (23)$$

IV. DEVICE SWITCHING

The desired evolution of the temperature distribution can be used to compute the necessary control actions of individual appliances. This section identifies such control actions using the three types of switching events identified in Section II-C. These are computed as a function of the time of evaluation t_i , the compressor state $c_i \in \{0, 1\}$ during the preceding interval $(t_{i-1}, t_i]$, and the current device temperature T_i (assumed to be measured in the appliance at time t_i).

A. Forced switching

TCLs are forced to switch `on` or `off` when their temperatures exceed the permitted interval $[T_{\text{low}}(t), T_{\text{high}}(t)]$. From the linear scaling of the temperature distributions around the pivot temperature $R(t)$ with a factor $s(t)$, it follows that

$$T_{\text{low}}(t) = R(t) - (R(t) - T_{\min})s(t), \\ T_{\text{high}}(t) = R(t) - (R(t) - T_{\max})s(t).$$

At t_i , the refrigerator must act if these bounds are violated at the start of the next time interval:

$$[T_i \leq R_i^+ - (R_i^+ - T_{\min})s_i^+] \Rightarrow c_{i+1} := 0 \quad (24)$$

$$[T_i \geq R_i^+ - (R_i^+ - T_{\max})s_i^+] \Rightarrow c_{i+1} := 1 \quad (25)$$

B. Continuous-time switching

We now consider the continuous-time stochastic switching rates from `on` to `off` states ($r^{1 \rightarrow 0}(t)$) and vice versa ($r^{0 \rightarrow 1}(t)$), required to maintain the desired shape of the temperature distribution. The switching rates for the energy provision mode are defined in [7, Eqs. (48)-(52)]. Here, we generalise these expressions to cover both energy provision and absorption modes ($R(t) \in \{T_{\min}, T_{\max}\}$) and simplify them using the z -coordinate transformation. Finally, we specialise the expressions for trapezoidal integration with piecewise constant power references.

The derivative of β can be simplified by substitution using (12) and (15), resulting in:

$$\frac{d\beta(t; R)}{dt} = \frac{1}{z(t) - \zeta(t)} \frac{d\Pi(t)}{dt} + \alpha\beta(t; R) \frac{1 + \zeta(t) - \Pi(t)}{z(t) - \zeta(t)} \\ = \frac{1}{z(t) - \zeta(t)} \frac{d\Pi(t)}{dt} - \alpha\beta(t; R)(1 + \beta(t; R))$$

This substitution can be used in [7, Eq. (51)] (adjusted for general R). Further simplification follows from setting $d\Pi(t)/dt = 0$ (because we consider piecewise constant sections between t_i). We compute the intermediate quantity $\Xi(t)$, using the identity found in [7, Eq. (38)], again taking left and right limits due to discontinuity at t_i .

$$\Xi_i^\pm = \lim_{\varepsilon \downarrow 0} \Xi(T_i, t_i \pm \varepsilon) \\ = \alpha^2 \left(\frac{P_i^\pm + Q_i^\pm}{P_i^\pm Q_i^\pm} \right) (X_i^\pm Y_i^\pm) - \alpha^2 (1 + \beta_i^\pm) (X_i^\pm + Y_i^\pm) \quad (26a)$$

with

$$P_i^\pm = (T_i - T_{\text{off}}) + (T_{\text{off}} - R_i^\pm)(1 - s_i^\pm) \quad (26b)$$

$$Q_i^\pm = (T_i - T_{\text{on}}) + (T_{\text{on}} - R_i^\pm)(1 - s_i^\pm) \quad (26c)$$

$$X_i^\pm = (T_i - T_{\text{off}}) + (T_i - R_i^\pm)\beta_i^\pm \quad (26d)$$

$$Y_i^\pm = (T_i - T_{\text{on}}) + (T_i - R_i^\pm)\beta_i^\pm \quad (26e)$$

The stochastic transition rates at the left and right limits to t_i are computed from Ξ_i^\pm using [7, Eq. (49)] (adjusted for general R), resulting in

$$r_{i,\pm}^{1 \rightarrow 0} = \max\left(0, \frac{-\Xi_i^\pm}{\alpha(T_i - T_{\text{off}}) + \alpha\beta_i^\pm(T_i - R_i^\pm)}\right) \quad (27a)$$

$$r_{i,\pm}^{0 \rightarrow 1} = \max\left(0, \frac{-\Xi_i^\pm}{\alpha(T_i - T_{\text{on}}) + \alpha\beta_i^\pm(T_i - R_i^\pm)}\right) \quad (27b)$$

Midpoint integration between adjacent time instants t_{i-1} and t_i is used to determine the resulting switching probabilities, where switching is implemented at $t = t_i$:

$$\text{Pr}_{\text{cont},i}^{1 \rightarrow 0} = \frac{1}{2}\Delta t_i(r_{i-1,+}^{1 \rightarrow 0} + r_{i,-}^{1 \rightarrow 0}) \quad (28a)$$

$$\text{Pr}_{\text{cont},i}^{0 \rightarrow 1} = \frac{1}{2}\Delta t_i(r_{i-1,+}^{0 \rightarrow 1} + r_{i,-}^{0 \rightarrow 1}) \quad (28b)$$

Note that the rates at both ‘inner’ edges of the interval Δt_i are used: the ‘+’ side at t_{i-1} and the ‘-’ side at t_i .

C. Instantaneous switching

Finally, consider the instantaneous stochastic switching at time t_i due to mode changes (energy absorption, energy delivery) or changes in $\Pi(t)$. This results in a discontinuous change in the net heating rate $v(T, t)$, which can only be achieved by a fraction of devices switching on or off at t_i . Following [8], we compute the probability of switching from the `on` to `off` state at time t_i , for a refrigerator that is currently `on`, as

$$\text{Pr}_{\text{inst},i}^{1 \rightarrow 0} = \max\left(0, 1 - \frac{(T_i - T_{\text{off}}) + (T_i - R_i^+)\beta_i^+}{(T_i - T_{\text{off}}) + (T_i - R_i^-)\beta_i^-}\right) \quad (29a)$$

Note that the switching probability includes both a contribution from the discrete change in power level at t_i as well as a possible mode transition in the previous interval that is implemented at t_i . The switching probability for fridges in the `off` state, $c_i = 0$, is defined analogously as

$$\text{Pr}_{\text{inst},i}^{0 \rightarrow 1} = \max\left(0, 1 - \frac{(T_i - T_{\text{on}}) + (T_i - R_i^+)\beta_i^+}{(T_i - T_{\text{on}}) + (T_i - R_i^-)\beta_i^-}\right) \quad (29b)$$

D. Combined stochastic switching

Formally, the continuous-time (28) and instantaneous (29) switching probabilities should be evaluated in sequence, because the former occurs during the interval $(t_{i-1}, t_i]$ and the latter at time t_i . This would account for the possibility that an appliance switches off and on again within a single interval, or vice versa. Here, we assume that the switching probability associated with the continuous-time process is small to allow us to evaluate both probabilities in a single step.

$$\text{Pr}_i^{1 \rightarrow 0} = \text{Pr}_{\text{cont},i}^{1 \rightarrow 0} + \text{Pr}_{\text{inst},i}^{1 \rightarrow 0} \quad (30a)$$

$$\text{Pr}_i^{0 \rightarrow 1} = \text{Pr}_{\text{cont},i}^{0 \rightarrow 1} + \text{Pr}_{\text{inst},i}^{0 \rightarrow 1} \quad (30b)$$

V. ALGORITHM AND RESULTS

The discrete time algorithm for updating the compressor state derived in sections III and IV, is summarised in pseudocode in Algorithm 1. The algorithm was implemented in Python 3.7 using the `numba` package to benefit from just-in-time compilation for considerable speedups.

Algorithm 1 State update algorithm

```

function UPDATE_COMPRESSOR_STATE( $\Pi_{i+1}, T_i, t_i$ )
  # load previously computed information
  load appliance model  $\mathcal{M}$  and operating range  $w$ 
  load  $c_i, \Pi_i, z_{i-1}, t_{i-1}, r_{i-1,+}^{1 \rightarrow 0}, r_{i-1,+}^{0 \rightarrow 1}$ 
5:  # implement power and energy limits
  compute  $z_i$  using (17)
  if  $z_i \leq 0$  then                                     ▷ energy delivery mode
    if  $z_i \leq w\zeta(T_{\text{max}})$  then
       $\Pi_{i+1} \leftarrow \max(\Pi_{i+1}, 1 + w\zeta(T_{\text{max}}))$ 
10:  clip  $\Pi_{i+1}$  to limits in (22)
  else                                                 ▷ energy absorption mode
    if  $z_i \geq w\zeta(T_{\text{min}})$  then
       $\Pi_{i+1} \leftarrow \min(\Pi_{i+1}, 1 + w\zeta(T_{\text{min}}))$ 
    clip  $\Pi_{i+1}$  to limits in (23)
15:  # determine distribution and switching variables
  compute  $R_i^\pm, \zeta_i^\pm, \beta_i^\pm, s_i^\pm$  using (18)-(20)
  compute  $r_{i,\pm}^{1 \rightarrow 0}, r_{i,\pm}^{0 \rightarrow 1}$  using (26)-(27)
  compute  $\text{Pr}_i^{1 \rightarrow 0}, \text{Pr}_i^{0 \rightarrow 1}$  using (28)-(30)
  # implement compressor switching
20:  if  $c_i = 1$  then                                     ▷ currently on
    if  $T_i \leq R_i^+ - (R_i^+ - T_{\text{min}})s_i^+$  then  $c_{i+1} \leftarrow 0$ 
    else
       $U \leftarrow$  uniform random  $\in [0, 1]$ 
      if  $U \leq \text{Pr}_i^{1 \rightarrow 0}$  then  $c_{i+1} \leftarrow 0$ 
25:  else  $c_{i+1} \leftarrow 1$                                ▷ remain on
    else                                                 ▷ currently off
      if  $T_i \geq R_i^+ - (R_i^+ - T_{\text{max}})s_i^+$  then  $c_{i+1} \leftarrow 1$ 
      else
         $U \leftarrow$  uniform random  $\in [0, 1]$ 
30:  if  $U \leq \text{Pr}_i^{0 \rightarrow 1}$  then  $c_{i+1} \leftarrow 1$ 
        else  $c_{i+1} \leftarrow 0$                                ▷ remain off
  return  $c_{i+1}$                                        ▷ updated compressor state

```

For simulations, thermal model parameters were taken from [8, domestic refrigerator class]: $\alpha = 1/7200\text{s}$; $T_{\text{max}} = 7^\circ\text{C}$; $T_{\text{min}} = 2^\circ\text{C}$; $T_{\text{on}} = -44^\circ\text{C}$; $T_{\text{off}} = 20^\circ\text{C}$. Heterogeneous appliances were generated from these parameters by individually multiplying them with a random factor that was uniformly distributed between 0.8 and 1.2. All appliances had a maximum power consumption $P_{\text{on}}^a = 70\text{W}$ and operating range $w = 0.9$ (not binding for the parameters used). Each appliance was randomly initialised as follows. The compressor was set to the `on` state with a probability equal to the steady state duty cycle P_0^a/P_{on}^a and the temperature was initialised according to the steady state probability distributions $f_0(T|c_0^a = 1) \propto 1/(T - T_{\text{on}})$ and $f_0(T|c_0^a = 0) \propto 1/(T_{\text{off}} - T)$.




Fig. 1. Reference signal (top) and response of a single appliance (middle and bottom).




Fig. 2. Response of a heterogeneous aggregate of appliances (top) and deviation from the reference (bottom).

Figure 1 shows a reference signal (top), with a length of 5 hours, that demonstrates a variety of features. The middle and bottom panels show the compressor state c_i and temperature T_i , respectively, of a single appliance that tracks the reference signal. These results illustrate the apparently weak relation between the reference signal and single device dynamics (middle), and the ability of the controller to strictly respect the temperature bounds (bottom).

Next, heterogeneous populations of 1,000 and 100,000

appliances were simulated, tracking the same reference signal. Figure 2 illustrates the convergence of the aggregate response to the reference signal as the number of independent appliances increases. The top panel shows absolute power consumption per appliance; the bottom panel the deviation from the reference.

VI. CONCLUSIONS AND FUTURE WORK

This paper has derived a discrete time TCL controller for decentralised demand response. The results illustrate the ability to accurately track a reference signal with a large population of heterogeneous appliances. Moreover, Algorithm 1 has low computational complexity, which permits implementation on embedded hardware with severe computational constraints, or it can be used to achieve efficient simulations. The simulation of 100,000 devices for 5 hours using 10s time steps took only 36 seconds (using an Intel i5-7360U CPU under macOS 10.14.2). Moreover, the ability to use variable time steps can further alleviate real time constraints.

Lab testing of the algorithm in a modified refrigerator is currently ongoing. Both in the lab and using simulations, it is of interest to investigate the robustness of the controller against perturbations from the idealised setting [10]. Initial simulation experiments suggest that the performance of the controller is quite robust to misspecification of the thermal model. Nevertheless, it is interesting to consider how the controller could be enhanced with a means for an appliance to learn and test its own thermal model.

REFERENCES

- [1] D. S. Callaway and I. A. Hiskens, "Achieving Controllability of Electric Loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, jan 2011.
- [2] D. Angeli and P.-A. Kountouriotis, "A Stochastic Approach to "Dynamic-Demand" Refrigerator Control," *IEEE Transactions on Control Systems Technology*, vol. 20, no. 3, pp. 581–592, may 2012.
- [3] S. P. Meyn, P. Barooah, A. Busic, Y. Chen, and J. Ehren, "Ancillary Service to the Grid Using Intelligent Deferrable Loads," *IEEE Transactions on Automatic Control*, vol. 60, no. 11, pp. 2847–2862, nov 2015.
- [4] L. C. Totu, R. Wisniewski, and J. Leth, "Demand Response of a TCL Population Using Switching-Rate Actuation," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 5, pp. 1537–1551, sep 2017.
- [5] S. Iacovella, P. Vingerhoets, G. Deconinck, N. Honeth, and L. Nordstrom, "Multi-Agent Platform for Grid and Communication Impact Analysis of Rapidly Deployed Demand Response Algorithms," in *EnergyCon 2016*, Leuven, 2016.
- [6] D. Ponce de Leon Barido, S. Suffian, J. Rosa, E. Brewer, and D. M. Kamen, "Enabling Micro-level Demand-Side Grid Flexibility in Resource Constrained Environments," in *2017 IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation*, Pittsburgh, PA, USA, 2017.
- [7] S. H. Tindemans, V. Trovato, and G. Strbac, "Decentralized Control of Thermostatic Loads for Flexible Demand Response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, sep 2015.
- [8] V. Trovato, S. H. Tindemans, and G. Strbac, "Leaky storage model for optimal multi-service allocation of thermostatic loads," *IET Generation, Transmission & Distribution*, vol. 10, no. 3, pp. 585–593, feb 2016.
- [9] S. H. Tindemans, V. Trovato, and G. Strbac, "Frequency control using thermal loads under the proposed ENTSO-E Demand Connection Code," in *2015 IEEE Eindhoven PowerTech*. IEEE, jun 2015, pp. 1–6.
- [10] E. C. Kara, M. Berges, and G. Hug, "Impact of Disturbances on Modeling of Thermostatically Controlled Loads for Demand Response," *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 1–1, 2015.