Low-complexity decentralized algorithm for aggregate load control of thermostatic loads

Tindemans, Simon H.; Strbac, Goran

DOI

[10.1109/TIA.2020.3034889](https://doi.org/10.1109/TIA.2020.3034889)

Publication date

2021

Document Version

Final published version

Published in

IEEE Transactions on Industry Applications

Citation (APA)

Tindemans, S. H., & Strbac, G. (2021). Low-complexity decentralized algorithm for aggregate load control of thermostatic loads. *IEEE Transactions on Industry Applications*, 57(1), 987 - 998. [9246256].
<https://doi.org/10.1109/TIA.2020.3034889>

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 Decentralized Algorithm for Aggregate Load Control of Thermostatic Loads

Simon H. Tindemans , *Member, IEEE*, and Goran Strbac , *Member, IEEE*

Abstract—Thermostatically controlled loads such as refrigerators are exceptionally suitable as a flexible demand resource. This article derives a decentralized 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. Tracking accuracy is quantified as a function of population size and time step size, and correlations in the tracking error are investigated. The controller is shown to be robust to errors in model specification and to sudden perturbations in the form of random refrigerator door openings.

Index Terms—Aggregation, decentralized control, demand response, embedded controller, thermostatically controlled loads.

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].

The fundamental challenge in deploying TCLs as a flexibility resource stems from the fact that most such loads only have high-power and low-power states (although exceptions exist, and lead to distinct challenges [2]). Although the concept of frequency response from such on/off TCLs has been around for decades [3], it has gained popularity in recent years, since Short *et al.* [4] illustrated how a simple adaptation of a hysteresis controller could be used to contribute to frequency stability.

Manuscript received June 8, 2020; revised September 11, 2020; accepted October 17, 2020. Date of publication October 30, 2020; date of current version December 31, 2020. Paper 2020-MCRE-0336.R1, approved for publication in the IEEE TRANSACTIONS ON INDUSTRY APPLICATIONS by the Electric Machines Committee of the IEEE Industry Application Society. This work was supported by EDF Energy R&D UK Centre as part of the STAMINA project. (Corresponding author: Simon H. Tindemans.)

Simon H. Tindemans is with the Department of Electrical Sustainable Energy, Delft University of Technology, 2628 Delft, The Netherlands (e-mail: s.h.tindemans@tudelft.nl).

Goran Strbac is with the Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2BU, U.K. (e-mail: g.strbac@imperial.ac.uk).

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIA.2020.3034889

The simple control mechanism is effective, but may cause synchronization of appliances that are modulated by a common frequency signal [5]. Ad hoc randomization strategies have been proposed to prevent this problem [6], but recently Webborn [7] has argued that the inherent heterogeneity of appliances is sufficient to dampen this effect in practical applications.

A second development has been a desire to develop controllers that cannot only stabilize the grid, but accurately track regulation signals. This can be achieved in a (conceptually) straightforward manner using a real-time communication infrastructure for direct measurement and control of appliances (e.g., [8]). However, considerations of privacy and fault-tolerance mean that it is generally preferred to consider systems that are only reliant on global or infrequent measurements. This can be achieved by incorporating models of device-level temperature dynamics [9], [10] and considering aggregate population dynamics, for example in the form of Markov transition models (“temperature bin models”) [11] that are used to improve aggregate state estimation.

Moreover, since many TCLs are deployed in large numbers, control of their aggregate power is a natural application for randomized control schemes. An early such scheme [5] has the drawbacks that it controls only asymptotic power consumption and cannot safeguard the temperature constraints of individual appliances. More recent articles employing randomized control [12]–[16] use an explicit representation of temperatures to prevent unwanted temperature excursions. Within these, [13], [14] focus on realism, incorporating details such as limits on compressor cycling and nonconstant power demand. On the other hand, [16] focuses on the development of a more general continuous-time control framework.

This article extends the decentralized randomized control strategy that was introduced in [12], [17]. That strategy has the desirable feature that it requires only one-way broadcast information, yet achieves tracking of a reference signal that is immediate and 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. The importance of such performance considerations for practical implementations can be seen in [18] and [19]. Against this background, the contributions of the present article are as follows.

- 1) The distribution-referred approach to aggregate load control is concisely described in Section II and the associated equations are restated in natural coordinates in Section III.

- 2) A discrete-time controller is derived that is designed to track a piecewise-constant reference signal. It consists of two complementary parts: *distribution control* (Section III) and *device switching* (Section IV). The algorithm is particularly 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.
- 3) An explicit control algorithm is stated (Section V) and an open source Python implementation is made available.
- 4) Simulations confirm the ability of the proposed algorithm to track a (potentially discontinuous) reference signal with a heterogeneous population of TCLs. The tracking error reduces with fleet size.
- 5) The tracking error is quantified as a function of time step size, and correlations in the tracking error are investigated.
- 6) The controller is demonstrated to be robust to errors in model specification and to sudden state perturbations in the form of random refrigerator door openings.

This article is an extended version of the conference paper [20], which first introduced the discrete-time formulation of the controller and demonstrated its tracking performance for large fleets (contributions 1–4). Compared to the earlier article, this article includes an extended discussion of the probabilistic formalism underlying the controller, a functionally equivalent but more compact statement of the control algorithm, and additional analysis of the controller performance (contributions 5–6).

II. PRELIMINARIES

A. Appliance Model

Throughout this article, we consider a first-order TCL model [9], [10], expressed by the following differential equation [5] for the temperature T^a of the compartment of appliance a :

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

Here, $c^a(t) \in \{0, 1\} \equiv \{\text{off}, \text{on}\}$ 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. 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. For the refrigeration case, the model parameters are related to physical parameters as follows [13]: $\alpha^a = 1/(R^a C^a)$, where C^a is the thermal capacitance and R^a is the thermal resistance. T_{off}^a is the ambient temperature and $T_{\text{on}}^a = T_{\text{off}}^a - \eta^a R^a P_{\text{ON}}^a$, with η^a the coefficient of performance and P_{ON}^a the work performed by the compressor. 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 the 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 .

We can transition to a probabilistic framing of the control approach in the following manner. Let us consider that the appliance a has a known model $\mathcal{M}^a = (\alpha^a, P_{\text{ON}}^a, T_{\text{off}}^a, T_{\text{on}}^a, T_{\text{min}}^a, T_{\text{max}}^a)$. A *hypothetical* large population of independent devices with identical models $\mathcal{M}^a$ would have internal states $(T^a(t), c^a(t))$ that are distributed according to a steady-state distribution $\mathcal{D}_0^a(\mathcal{M}^a)$. Now consider a *single* appliance a , with an unknown (to the external observer) internal state $S^a = (T^a(t), c^a(t))$, i.e., an unknown initial condition. We can then say that its state is a random variable that is drawn from the steady-state distribution, $S^a \sim \mathcal{D}_0^a(\mathcal{M}^a)$. Moreover, if appliances are independently operated, this state is independent of the states of other appliances (regardless of whether their models are different). Then, the power consumption of device a is in effect a random process, with the (constant) expectation

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

where [12, eq. (33)]

$$P_0^a = \left[\frac{\log \left(\frac{T_{\text{max}}^a - T_{\text{on}}^a}{T_{\text{min}}^a - T_{\text{on}}^a} \right)}{\log \left(\frac{(T_{\text{max}}^a - T_{\text{on}}^a)(T_{\text{min}}^a - T_{\text{off}}^a)}{(T_{\text{min}}^a - T_{\text{on}}^a)(T_{\text{max}}^a - T_{\text{off}}^a)} \right)} \right] P_{\text{ON}}^a. \quad (4)$$

The expectation (3) is therefore taken over the internal states that an appliance might have at time t , given its model $\mathcal{M}^a$. Because no external control is applied and the initial state of the appliance was unknown, the result does not depend on time.

B. Aggregate Power Modulation

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

$$P_{\Pi}(t) = \sum_{a \in \mathcal{A}} P_{\Pi}^a(t) \quad (5)$$

of a collection of appliances $\mathcal{A}$. In this article, we consider the control approach introduced in [12], which modulates the power consumption using a broadcast signal $\Pi(t)$. This *reference signal* is available to all devices, and indicates the desired power consumption relative to the nominal power consumption of the device. Each device individually adapts its power consumption $P_{\Pi}^a(t)$ in order to satisfy

$$E_{S^a(t)}[P_{\Pi}^a(t)] = \Pi(t) P_0^a. \quad (6)$$

By design, a signal $\Pi(t) = 1$ results in the steady state power consumption (3), and changes in $\Pi(t)$ are immediately reflected in the expected power consumption. Moreover, the controller maintains independence between appliances: the random states $S^a(t)$ and $S^b(t)$ are independent of each other, when conditioned on the control signal $\Pi(\cdot)$. It thus follows from the Lyapunov condition (Lyapunov central limit theorem) that the aggregate power consumption (5) is given by

$$P_{\Pi}(t) = \Pi(t) \sum_{a \in \mathcal{A}} P_0^a + O(\sqrt{|\mathcal{A}|}) \quad (7)$$

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 [12]. In [17], the control signal was generated using a mixture of offline scheduling and real-time control, and in [21] various frequency-sensitive controllers to locally compute $\Pi(t)$ (e.g., a simple droop controller) were implemented.

C. Distribution-Referred Control

The controller introduced in [12] and [17] can be thought of as *distribution-referred*, because a device first derives a generic control model for devices of *its type*, by considering its thermal model $\mathcal{M}^a$ without reference to its actual state $(T^a(t), c^a(t))$. In accordance with the probabilistic frame defined in the previous section, each device considers a probability distribution of temperatures for appliances with its model $\mathcal{M}^a$ and an unknown state $\mathcal{S}^a$. One convenient approach is to parameterize this temperature distribution as a family $f_z(T; \mathcal{M}^a)$ that varies continuously in the parameter z , containing as a special case the steady-state temperature distribution [12, eqs. (31) and (32)]

$$f_0(T; \mathcal{M}^a) = \frac{k^a}{(T_{\text{off}}^a - T)(T - T_{\text{on}}^a)} \quad (8)$$

with

$$k^a = \frac{T_{\text{off}}^a - T_{\text{on}}^a}{\log \left(\frac{(T_{\text{max}}^a - T_{\text{on}}^a)(T_{\text{min}}^a - T_{\text{off}}^a)}{(T_{\text{min}}^a - T_{\text{on}}^a)(T_{\text{max}}^a - T_{\text{off}}^a)} \right)}. \quad (9)$$

At a high level, the controller for each appliance consists of two parts, which are evaluated in order:

- 1) *Distribution control*. Control the evolution of the distribution parameter $z(t)$ in such a way that the power consumption tracks the reference signal $\Pi(t)$ according to (6). 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) *Device switching*. Based on the *actual* appliance state $(T^a(t), c^a(t))$, compute stochastic control actions, in the form of ON/OFF switching. Switching events can be initiated in three distinct ways.
 - a) Forced switching (deterministic) when temperature limits $T_{\text{min}}(t)$ or $T_{\text{max}}(t)$ are exceeded.
 - b) A continuous-time switching process (stochastic) at intermediate temperatures in order to shape the temperature distribution.
 - c) Instantaneous switching (stochastic) on discrete changes of power setpoints, or when the controller switches between energy-provision and energy-absorption modes [17] (see also Section III-B).

The distribution control and device 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. Time-Discretization Procedure

The derivation of the controller in [12] was performed in continuous time and temperature coordinates. However, implementation in device controllers is greatly simplified when a discrete-time formulation can be given, which is the main contribution of this article. To move from a continuous-time to a discrete-time formulation, we partition the timeline by the ordered sequence of times $\{t_i\}$, indexed by the integer i , at which the controller is invoked. These 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

$$\Pi(t) = \Pi_i, \quad \text{for } t \in (t_{i-1}, t_i]. \quad (10)$$

The controller thus receives at t_i a new reference power level Π_{i+1} that must be applied for the upcoming interval $(t_i, t_{i+1}]$. Note that t_{i+1} is not necessarily known at time t_i .

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 discretized approximation of the continuous time controller, they will be synchronized with the control execution times t_i as follows. It is assumed that switching is immediate (at t_i).

- 1) A violation of the temperature limits will trigger corrective switching as soon as it is detected.
- 2) Switching between energy-provision and energy-absorption modes also results in a discontinuous change in the desired local heating/cooling rate. This is achieved by ON/OFF switching, immediately when a change in regime is detected.
- 3) 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. CONTROLLER: MANIPULATION OF DISTRIBUTION

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 $\Pi(t) = \Pi_i$, for $t \in (t_{i-1}, t_i]$. 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 [12, eq. (26)]. With the convention that $t_{-1} = -\infty$ and $\Pi_0 = 1$ (assuming an initial steady state), it follows that

$$\bar{T}(t_i) = T_{\text{off}} - \alpha(T_{\text{off}} - \bar{T}_0) \sum_{j=0}^i \Pi_j \int_{t_{j-1}}^{t_j} e^{-\alpha(t_i - t')} dt' \quad (11)$$

where the steady state average temperature $\bar{T}_0$ is computed using [12, eqs. (23) and (32) and (33)] as

$$\bar{T}_0 = T_{\text{off}} - k \times \log \left(\frac{T_{\text{max}} - T_{\text{on}}}{T_{\text{min}} - T_{\text{on}}} \right) \quad (12)$$

with k defined in (9). We define the dimensionless variable

$$z(t) = \frac{\bar{T}_0 - \bar{T}(t)}{T_{\text{off}} - \bar{T}_0} \quad (13)$$

to parameterize the distributions $f_{z(t)}(T)$, and simplify the notation in what follows. z is a representation of the cooling energy stored in the aggregate device population, relative to the uncontrolled state ($z = 0$).¹

B. Controller Modes

The algorithm in [12] implicitly generates the family of temperature distributions $f_{z(t)}(T)$ by the net heating rate $v(T, t)$, which is determined by averaging the heating/cooling rate over devices in the *off* (heating) state and *on* (cooling) state at time t and temperature T . The net heating rate and $f_{z(t)}$ are coupled through the continuity equation

$$\frac{\partial}{\partial t} f_{z(t)}(T) = -\frac{\partial}{\partial T} [v(T, t) f_{z(t)}(T)]. \quad (14)$$

The heating rate can be chosen freely between the lower and upper limits that are defined by the (negative) heating rate of devices in the *on* state ($v_{\text{on}}(T) = -\alpha(T - T_{\text{on}})$) and the heating rate of devices in the *off* state ($v_{\text{off}}(T) = \alpha(T_{\text{off}} - T)$). A convenient choice is the net heating rate

$$v(T, t) = \alpha\beta(t)(T - T_{\text{max}}) \quad (15)$$

that is parameterized by a single *control parameter* $\beta(t)$. 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 [17], it was coupled to a “mirrored” controller that is capable of absorbing energy from the grid by contracting to the pivot temperature T_{min} . The controller switches between *energy absorption* and *energy provision* modes whenever $\bar{T}(t)$ crosses $\bar{T}_0$ (when $z(t)$ crosses 0).

A generalized 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:

$$R(t) = \begin{cases} T_{\text{max}}, & \text{if } \bar{T}(t) \geq \bar{T}_0 \\ T_{\text{min}}, & \text{if } \bar{T}(t) < \bar{T}_0 \end{cases} \quad (16)$$

C. Control Parameter

The control parameter $\beta(t; R)$ is determined by the desired reference power $\Pi(t)$ according to [12, eq. (36)]:

$$\begin{aligned} \beta(t; R) &= \frac{\Pi(t)(T_{\text{off}} - \bar{T}_0) - (T_{\text{off}} - \bar{T}(t))}{R(t) - \bar{T}(t)} \\ &= \frac{(\Pi(t) - 1) - z(t)}{z(t) - \zeta(R(t))} \end{aligned} \quad (17)$$

where

$$\zeta(R) = \frac{\bar{T}_0 - R}{T_{\text{off}} - \bar{T}_0}. \quad (18)$$

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} [12]. The extent of this scaling at time t_i is compactly represented by the scale parameter

$$\begin{aligned} s(t) &= \frac{R(t) - \bar{T}(t)}{R(t) - \bar{T}_0} \\ &= 1 - z(t)/\zeta(R(t)). \end{aligned} \quad (19)$$

E. Discretization

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 discretization $z_i = z(t_i)$, which is computed from (11) as

$$z_i = \sum_{j=0}^i (\Pi_j - 1) \left(e^{-\alpha(t_i - t_j)} - e^{-\alpha(t_i - t_{j-1})} \right). \quad (20)$$

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

$$z_i = z_{i-1} e^{-\alpha\Delta t_i} + (\Pi_i - 1)(1 - e^{-\alpha\Delta t_i}). \quad (21)$$

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

$$\hat{R}(t) = R_i, \quad \text{for } t \in (t_{i-1}, t_i]. \quad (22)$$

¹Note that z is related to the variable σ used in [17] as $z = \sigma - 1$.

with

$$R_{i+1} = \begin{cases} T_{\max}, & \text{if } z_i \leq 0 \\ T_{\min}, & \text{if } z_i > 0. \end{cases} \quad (23)$$

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

$$R_i^- = \lim_{\varepsilon \downarrow 0} \hat{R}(t - \varepsilon) = R_i \quad (24a)$$

$$R_i^+ = \lim_{\varepsilon \downarrow 0} \hat{R}(t + \varepsilon) = R_{i+1}. \quad (24b)$$

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

$$\zeta_i^\pm = \frac{\bar{T}_0 - R_i^\pm}{T_{\text{off}} - \bar{T}_0}, \quad s_i^\pm = 1 - z_i / \zeta_i^\pm \quad (25)$$

$$\beta_i^- = \frac{(\Pi_i - 1) - z_i}{z_i - \zeta_i^-}, \quad \beta_i^+ = \frac{(\Pi_{i+1} - 1) - z_i}{z_i - \zeta_i^+}. \quad (26)$$

F. Energy 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 $f_z(T)$ exist outside of this domain). However, operation near the limits is infeasible in practice, because appliances would need to switch at very high rates to operate in a narrow temperature range. Therefore, we shall use a restricted range of operating temperatures that is scaled with a fraction $w < 1$ around the steady state operating temperature $\bar{T}_0$

$$(1 - w)\bar{T}_0 + wT_{\min} \leq \bar{T}(t) \leq (1 - w)\bar{T}_0 + wT_{\max}. \quad (27)$$

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

$$w\zeta(T_{\max}) \leq z(t) \leq w\zeta(T_{\min}). \quad (28)$$

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 (21). This leads to the update rule for Π_{i+1} :

$$\Pi_{i+1} = \begin{cases} \max(\Pi_{i+1}, 1 + w\zeta(T_{\max})), & \text{if } z_i \leq w\zeta(T_{\max}) \\ \min(\Pi_{i+1}, 1 + w\zeta(T_{\min})), & \text{if } z_i \geq w\zeta(T_{\min}) \\ \Pi_{i+1}, & \text{otherwise} \end{cases} \quad (29)$$

Note that $\beta_i^\pm$ will need to be (re-)computed to ensure consistency with an adjusted power level.

G. Power Constraints

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. Following the procedure in [12], these limits can be derived by requiring that the desired *average* rate of temperature change for devices with temperature T_i remains within the physical limits imposed by the `on` and `off` states,

i.e., $v_{\text{ON}}(T_i) \leq v(T_i, t_i) \leq v_{\text{OFF}}(T_i)$ for all T_i and t_i . Expanding these for the pivoting controller at time $t_i + \varepsilon$ yields

$$-\alpha(T_i - T_{\text{on}}) \leq \alpha\beta_i^+(T_i - R_i^+) \leq -\alpha(T_i - T_{\text{off}}). \quad (30)$$

Since this equality is linear in T_i and must hold for all devices, it suffices to verify this constraint at the most extreme attainable temperatures. From the linear scaling of the temperature distributions around the pivot temperature $R(t)$ with a factor $s(t)$, it follows that the permitted temperature interval $[T_{\text{low}}(t), T_{\text{high}}(t)]$ is defined by

$$T_{\text{low}}(t) = R(t) - (R(t) - T_{\min})s(t) \quad (31a)$$

$$T_{\text{high}}(t) = R(t) - (R(t) - T_{\max})s(t). \quad (31b)$$

Depending on the controller mode, one of these limits is equal to the pivoting temperature $R(t)$. Since we have $v(R(t), t) = 0$ by definition, (30) is trivially satisfied at this limit. We can therefore focus on the opposite limit, which can be written as

$$T_{\text{limit}}(t) = R(t) + [T_{\min} + T_{\max} - 2R(t)]s(t). \quad (32)$$

Setting $T_i = T_{\text{limit}}(t_i + \varepsilon)$ in (30) and reordering of the result yields the following power constraints on Π_{i+1} :

$$\begin{aligned} 1 + \zeta_i^+ \left(\frac{T_{\min} + T_{\max} - T_{\text{off}} - R_i^+}{T_{\min} + T_{\max} - 2R_i^+} \right) &\leq \Pi_{i+1} \\ &\leq 1 + \zeta_i^+ \left(\frac{T_{\min} + T_{\max} - T_{\text{on}} - R_i^+}{T_{\min} + T_{\max} - 2R_i^+} \right). \end{aligned} \quad (33)$$

This result unifies the separately calculated limits for the energy-provision ($z_i < 0$) and energy-absorption ($z_i > 0$) regimes in [12], [17].

IV. CONTROLLER: 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

$$T_{\text{low}}(t) = R(t) - (R(t) - T_{\min})s(t) \quad (34a)$$

$$T_{\text{high}}(t) = R(t) - (R(t) - T_{\max})s(t). \quad (34b)$$

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

$$[T_i \leq R_i^+ - (R_i^+ - T_{\min})s_i^+] \Rightarrow c_{i+1} := 0 \quad (35a)$$

$$[T_i \geq R_i^+ - (R_i^+ - T_{\max})s_i^+] \Rightarrow c_{i+1} := 1. \quad (35b)$$

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 [12, eqs. (48)–(52)]. Here, we generalize 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 specialize the expressions for trapezoidal integration with piecewise constant power references.

The derivative of β can be simplified by substitution using (13) and (18), resulting in

$$\begin{aligned} \frac{d\beta(t; R)}{dt} &= \frac{1}{z(t) - \zeta(t)} \frac{d\Pi(t)}{dt} + \alpha\beta(t; R) \frac{1 + \zeta(t) - \Pi(t)}{z(t) - \zeta(t)} \\ &= \frac{1}{z(t) - \zeta(t)} \frac{d\Pi(t)}{dt} - \alpha\beta(t; R)(1 + \beta(t; R)). \end{aligned} \quad (36)$$

This substitution can be used in [12, 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 [12, eq. (38)], again taking left and right limits due to discontinuity at t_i .

$$\begin{aligned} \Xi_i^\pm &= \lim_{\varepsilon \downarrow 0} \Xi(T_i, t_i \pm \varepsilon) \\ &= \alpha^2 \left(\frac{P_i^\pm + Q_i^\pm}{P_i^\pm Q_i^\pm} \right) (X_i^\pm Y_i^\pm) - \alpha^2 (1 + \beta_i^\pm) (X_i^\pm + Y_i^\pm) \end{aligned} \quad (37a)$$

with

$$P_i^\pm = (R_i^\pm - T_{\text{off}})s_i^\pm + (T_i - R_i^\pm) \quad (37b)$$

$$Q_i^\pm = (R_i^\pm - T_{\text{on}})s_i^\pm + (T_i - R_i^\pm) \quad (37c)$$

$$X_i^\pm = (T_i - T_{\text{off}}) + (T_i - R_i^\pm)\beta_i^\pm \quad (37d)$$

$$Y_i^\pm = (T_i - T_{\text{on}}) + (T_i - R_i^\pm)\beta_i^\pm. \quad (37e)$$

The stochastic transition rates at the left and right limits to t_i are computed from $\Xi_i^\pm$ using [12, eq. (49)] (adjusted for general R), resulting in

$$r_{i,\pm}^{1 \rightarrow 0} = \max \left(0, -\frac{\Xi_i^\pm}{\alpha X_i^\pm} \right) \quad (38a)$$

$$r_{i,\pm}^{0 \rightarrow 1} = \max \left(0, -\frac{\Xi_i^\pm}{\alpha Y_i^\pm} \right). \quad (38b)$$

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 $t = t_i$

$$\text{Pr}_{\text{cont},i}^{1 \rightarrow 0} = \frac{1}{2} \Delta t_i (r_{i-1,+}^{1 \rightarrow 0} + r_{i,-}^{1 \rightarrow 0}) \quad (39a)$$

$$\text{Pr}_{\text{cont},i}^{0 \rightarrow 1} = \frac{1}{2} \Delta t_i (r_{i-1,+}^{0 \rightarrow 1} + r_{i,-}^{0 \rightarrow 1}). \quad (39b)$$

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 [17], we compute the probability of switching from the `on` to `off` state at time t_i , for a refrigerator that is currently `on`, as

$$\text{Pr}_{\text{inst},i}^{1 \rightarrow 0} = \max \left(0, 1 - \frac{X_i^+}{X_i^-} \right). \quad (40a)$$

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

$$\text{Pr}_{\text{inst},i}^{0 \rightarrow 1} = \max \left(0, 1 - \frac{Y_i^+}{Y_i^-} \right). \quad (40b)$$

D. Combined Stochastic Switching

Formally, the continuous-time (39) and instantaneous (40) 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.

$$\text{Pr}_i^{1 \rightarrow 0} = \text{Pr}_{\text{cont},i}^{1 \rightarrow 0} + \text{Pr}_{\text{inst},i}^{1 \rightarrow 0} \quad (41a)$$

$$\text{Pr}_i^{0 \rightarrow 1} = \text{Pr}_{\text{cont},i}^{0 \rightarrow 1} + \text{Pr}_{\text{inst},i}^{0 \rightarrow 1}. \quad (41b)$$

V. ALGORITHM AND IMPLEMENTATION

The discrete time algorithm for updating the compressor state derived in Sections III and IV is summarized in pseudocode in Algorithm 1. The algorithm was implemented in Python 3.7.7 using the `numba` package to benefit from just-in-time compilation for considerable speedups. The source code for the controller and all results in this article has been released under the MIT license [22]; an interactive compute capsule that reproduces all figures and numerical results is available at [23].

Fig. 1 provides a schematic description of the relation between the control algorithm (*local TCL controller* block) and the other components of the overall system. The behavior of each TCL is determined by the interaction between its physical device characteristics, its environment (e.g., ambient temperature and door openings) and the compressor signal c_i . The global TCL controller—which could be a transmission system operator or a demand response service provider—monitors the condition of the power system or its associated markets and sends the dispatch signal $\Pi(t_i)$ to all TCLs.

Fig. 1. Schematic depiction of the local and global control blocks.

The TCL controller makes use of this signal, in combination with the current time t_i and the measured cabinet temperature T_i . It further uses a persistent controller state, and a representation of the device model $\mathcal{M}$. Ideally, this model is a good representation of the physical device characteristics, including its environment and usage. In practice, there will be parameter and model errors to contend with, leading to unavoidable control errors. In advanced implementations, the model can be refined on the basis of observed temperature measurements.

Due to the conceptual separation between the state update control algorithm and the physical evolution of the system, the control algorithm is equally suitable for embedded applications as for simulations. When embedded in a real system, all other blocks except for the global TCL controller are physical. In a simulation environment, the other blocks are all provided by the modeler.

VI. RESULTS

For simulations, thermal model parameters were taken from [17, domestic refrigerator class]: $\alpha = 1/7200 \text{ s}^{-1}$; $T_{\max} = 7^\circ\text{C}$; $T_{\min} = 2^\circ\text{C}$; $T_{\text{on}} = -44^\circ\text{C}$; $T_{\text{off}} = 20^\circ\text{C}$. Heterogeneous appliances were generated from these parameters by individually multiplying their values with a random factor that was uniformly distributed between 0.8 and 1.2. The temperature evolution of each appliance was computed by integrating (1) using the Euler method; the integration time steps equaled those of the controller, unless otherwise noted. 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 initialized 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 initialized in the range $[T_{\min}, T_{\max}]$ according to the steady-state probability distributions $f_0(T|c_0^a = 1) \propto 1/(T - T_{\text{on}})$ and $f_0(T|c_0^a = 0) \propto 1/(T_{\text{off}} - T)$.

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:  # update population parameters
    compute  $z_i$  using (21)
    compute  $R_i^\pm, \zeta_i^\pm, s_i^\pm$  using (23)-(25)

10: # implement power and energy limits
    if  $z_i \leq w\zeta(T_{\max})$  then lower energy limit violated
       $\Pi_{i+1} \leftarrow \max(\Pi_{i+1}, 1 + w\zeta(T_{\max}))$ 
    if  $z_i \geq w\zeta(T_{\min})$  then upper energy limit violated
       $\Pi_{i+1} \leftarrow \min(\Pi_{i+1}, 1 + w\zeta(T_{\min}))$ 

15: clip  $\Pi_{i+1}$  to power limits in (33)
    compute  $\beta_i^\pm$  using (26)

    # determine distribution and switching variables
    compute  $r_{i,\pm}^{1 \rightarrow 0}, r_{i,\pm}^{0 \rightarrow 1}$  using (37)-(38)

20: compute  $\text{Pr}_i^{1 \rightarrow 0}, \text{Pr}_i^{0 \rightarrow 1}$  using (39)-(41)

    # implement compressor switching
    if  $c_i = 1$  then ▷ currently on
      if  $T_i \leq R_i^+ - (R_i^+ - T_{\min})s_i^+$  then  $c_{i+1} \leftarrow 0$ 
    25: else
       $u \leftarrow \text{uniform random} \in [0, 1]$ 
      if  $u \leq \text{Pr}_i^{1 \rightarrow 0}$  then  $c_{i+1} \leftarrow 0$ 
      else  $c_{i+1} \leftarrow 1$  ▷ remain on
    else ▷ currently off
    30: if  $T_i \geq R_i^+ - (R_i^+ - T_{\max})s_i^+$  then  $c_{i+1} \leftarrow 1$ 
      else
       $u \leftarrow \text{uniform random} \in [0, 1]$ 
      if  $u \leq \text{Pr}_i^{0 \rightarrow 1}$  then  $c_{i+1} \leftarrow 1$ 
      else  $c_{i+1} \leftarrow 0$  ▷ remain off

35: return  $c_{i+1}$  ▷ updated compressor state

```

A. Single Device Analysis

Fig. 2 shows a reference signal (top), with a length of 5 h, that demonstrates a variety of features. The second panel shows the resulting value of the population parameter z , indicating that the controller is predominantly in the energy-absorption mode ($z \geq 0$) during the first 3 h, before switching to the energy-provision mode ($z \leq 0$) for the remainder of the time.

The bottom two panels show the compressor state c_i and temperature T_i , respectively, of a single appliance that tracks the reference signal. The control state of a single appliance does not have a very apparent relation with the reference signal. The bottom panel indicates the device temperature T_i (solid line) alongside the lower and upper bound temperatures $T_{\text{low}}(t)$ and $T_{\text{high}}(t)$ (dashed lines). This illustrates the ability of the controller to strictly respect the temperature bounds. It also demonstrates that on/off switching events occur both due to the device reaching its temperature bounds (deterministic) and due to stochastic switching events at intermediate temperatures.

Fig. 2. Reference signal (top) and corresponding population parameter (second). The bottom two panels show the state evolution of a random single appliance (compressor state and temperature). 10 s time steps were used for the simulation.

B. Population Size

Next, heterogeneous populations of 1000 and 100 000 appliances were simulated, tracking the same reference signal. Fig. 3 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 in Watt per appliance.

To further analyze the convergence as a function of device counts, the signal was repeated five consecutive times to obtain a simulated 25-h run. The per-appliance power deviation for these runs amounted to 0.864 and 0.116 W, respectively, for populations of 1000 and 100 000 appliances. Although this reduction in standard deviation is substantial (7.4 $\times$), it nevertheless falls short of the 10 $\times$ reduction predicted by (7). This discrepancy can largely be attributed to the 10-s step size used for the simulations: devices that exceed temperature bounds or are affected by a stochastic switching event must wait until the next 10-s barrier to implement this event. Simulating 100 000 appliances with 1-s time steps (and a reference signal with stepwise increments at 10-s intervals, for a fair comparison) resulted in a standard deviation of only 0.098 W per appliance. This number can be compared to the total signal magnitude of approximately 10 W per appliance.

Fig. 3. Response of a heterogeneous aggregate of appliances (top) and deviation from the reference (bottom). 10 s time steps were used for the simulation.

Fig. 4. Deviation between the reference signal from Fig. 3 and response of 100 000 appliances for various time steps. The top panel compares fixed time steps of various sizes. The bottom panel compares a fixed 20 s time step with a 10 s time step where half of controller invocations is dropped at random.

C. Simulation Step Size

The impact of the time step size is further investigated in Fig. 4. The top panel compares the reference tracking accuracy for time steps of 1, 10, and 60 s. Whereas 60 s time steps result in significant “crosstalk” from the reference signal on the deviation signal, this is much reduced for 10 s time steps and nearly absent (at least visually) for 1 s time steps. This suggests that 10 s time steps are sufficient for all but the most critical applications (for this population size of 100 000 appliances).

The bottom panel compares the aggregate tracking accuracy obtained using fixed time steps of 20 s and variable time steps

Fig. 5. Autocorrelation results obtained with populations of 100 000 appliances and 25-h simulation runs. Top panels show results for signal tracking simulations; bottom panels for steady-state simulations. Solid lines indicate results from 10-s time steps and dashed lines using 1-s time steps. Values within shaded areas are nonsignificant at the 95% confidence level.

with an *average* duration of 20 s. Variable time steps were generated by running a simulation with 10 s time steps and skipping each time step with 50% probability—independently for each time step and device. Fig. 4 shows that the tracking performance is similar, with the exception of discontinuous transitions in the reference power, where some devices with variable time steps are unable to instantly track the desired change. However, the tracking fully recovers after a short delay. Moreover, the fixed time step controller *is* able to track all discontinuities in the reference signal instantaneously.

D. Temporal Correlations of Tracking Errors

In this section, we look closer at the temporal signature of the tracking errors using their autocorrelation, shown in Fig. 5. The panels on the right show that, for homogeneous populations of TCLs, long-range correlations are present in the deviation from the theoretical mean, even resulting in *exact* repetition in the steady state. However, when the population tracks a signal (top right), this long-range autocorrelation is reduced, because stochastic switching effectively randomizes the appliance states.

For heterogeneous populations, the autocorrelation signatures are much reduced, especially in the steady state. However, residual long-range correlations were still observed for the signal tracking case (top left; solid line), suggesting that they were induced by the control signal.

The solid lines are computed on the basis of 10 s time steps, which were shown in Section VI-C to be marginally sufficient for good tracking performance. Additional results using 1 s time steps (dashed lines in both top panels) show that this indeed reduces the long-term autocorrelations for the heterogeneous

Fig. 6. Comparison of aggregate tracking error of heterogeneous appliances for different types of model. Three cases are depicted: known models (red; also shown in Figs. 3 and 4), randomized model errors (blue), and the use of a common control model (green). Simulations used 100 000 devices and 10s time steps.

case, but not for the homogeneous population case. Moreover, the periodicity of both signals changes from that of the control signal (approx. 30 min) to that of the average cycle length of the appliances (approx. 1 h), in line with the steady-state results. This supports the hypothesis that a reduction in time step to 1 s is sufficient to eliminate the impact of the control signal for critical applications.

E. Model Error

The proposed control algorithm relies on a thermal model of the appliance to compute the appropriate switching parameters. In the previous sections, we have assumed that this model was known to the controller, but this is arguably a strong assumption that does not hold in practice [24]. The sensitivity to model errors is investigated in this section.

For this purpose, we distinguish the *physical model* that dictates the appliance physics according to (1) and the *control model* that is used in Algorithm 1. Three cases were considered.

- 1) Known models. The control model parameters were identical to those of the physical model.
- 2) Common control model. The physical model for each device was randomly drawn as described above, but devices share a nominal control model (with median parameter values). The latter can be thought of as a factory-embedded “archetype model” for the appliance type.
- 3) Randomized model errors. Parameters for the physical and control models were drawn independently. Such a scenario might occur if each appliance independently (and imperfectly) attempts to learn its own model.

In all cases, the control parameters $T_{\min}^a$ and $T_{\max}^a$ were taken from the physical model, as they represent the user’s cooling preferences, as adapted to the physical environment of the appliance. The physical model parameters were used to generate the initial random temperature of the device, which represents an initial state where a hysteresis controller was used prior to $t = 0$.

Fig. 6 compares the average tracking error for all three scenarios, when tracking the signal shown in Fig. 2 (top) with 100 000

Fig. 7. Results of an aggregate control experiment in the presence of door openings. Top: 24-h reference signal (a repeated test pattern) and the door opening rate with hourly changes. Middle: comparison of per-appliance power consumption with different baselines. Bottom: per-appliance tracking error between the controlled population and the control target.

devices. Model error is shown to slightly increase the tracking error, especially when following a discontinuous signal (2–3 h). However, the magnitude of this error remains well-controlled and the error signal quickly reduces to normal levels when the reference signal reverts to the nominal value $\Pi(t) = 1$. We emphasize that model parameters may be off by as much as 50% in the case of random model error, further illustrating the robustness of the controller design. However, we point out that a significant bias in the parameter estimation errors may lead to a deterioration of performance, as errors are less likely to cancel out.

F. Door Openings

The model errors analyzed in the previous section represent a constant discrepancy between the assumed thermal model and reality. In addition, it is worthwhile to consider *dynamic* disturbances such as the effect of changing outdoor temperatures [25] and the impact of frequent door openings [13]. In the following, we consider only the case of door openings, which induce sudden large changes in the model state (temperature). As such, they can be thought of as a severe stress test of the controller.

We model door openings to occur independently with a time-varying rate, adapted from [13, Fig. 15] and scaled to result in an average of 20 openings per day (see Fig. 7, top panel). Each door opening event lasts 20 s and results in a reduction of the thermal conductance R^a by a factor of 25 [13] while the door is open. For the simulation of temperature dynamics, the time step of Euler integration was reduced by the same factor of 25 to reduce numerical errors, but the time step for the controller remained the same at 10 s.

The results of this simulated experiment with 100 000 devices are shown in Fig. 7. The middle panel shows the expected power consumption (per appliance) in absence of door openings (dashed black line), alongside the simulated power consumption with door openings (solid black line). The first observation is that the addition of door openings leads to a significant and time-varying increase in average power consumption, by 11% on average and 44% at its peak.

The red line indicates the power consumption of appliances with door openings that simultaneously track the reference signal shown in Fig. 7 (top). In order to quantify the tracking performance, this should be compared with a target power consumption, for which the reference signal $\Pi(t)$ would be designed. The time-varying door openings result in a time-varying baseline for the power consumption, so the target power consumption can be given by

$$\bar{P}_{\Pi, \text{door}}^{\text{target}}(t) = \bar{P}_{0, \text{door}}(t) + (\Pi(t) - 1)\bar{P}_0 \quad (42)$$

where all powers $\bar{P}_x$ are taken to be averages over the population (i.e., W per device). The second term specifies the modulation of the power consumption with respect to the known steady-state power level $\bar{P}_0$. The baseline power consumption $\bar{P}_{0, \text{door}}(t)$ is variable, but it can be predicted in the long term. As no analytical expression is available for this baseline, the empirical power consumption (solid black line) for an uncontrolled population is used instead. The resulting target power consumption $\bar{P}_{\Pi, \text{door}}^{\text{target}}(t)$ is indicated by the blue line in the middle panel of Fig. 7. The bottom panel depicts the difference between this reference power consumption and the realized power consumption. It demonstrates that the controller is able to track the signal to within an error of ± 2 W/appliance, despite very large perturbations in

the form of door openings occurring at the same time as large swings in reference power consumption.

VII. CONCLUSION AND FUTURE WORK

This article has derived a discrete-time TCL controller for decentralized demand response. The results illustrate the ability to accurately track a (potentially discontinuous) reference signal using a large population of heterogeneous appliances. The tracking performance remained high even in the presence of significant model errors and random refrigerator door openings that cause unanticipated spikes in appliance temperatures.

Moreover, from a computational perspective, it was shown that the time-discretization procedure used led to acceptable tracking errors for time steps of 10 s, and further suppression of correlated errors with a smaller time step of 1 s. Coupled with the low computation complexity of Algorithm 1, this permits implementation on embedded hardware with severe computational constraints, or it can be used to achieve efficient offline simulations. The simulation of 100 000 devices for 5 h using 10 s time steps took only 35 s (using an Intel i5-7360 U CPU under macOS 10.15.4 and Python 3.7.7). Moreover, the ability to use variable time steps can further alleviate real time constraints.

In addition to these computational benefits, the controller inherits desirable properties from its continuous-time version [12] that enable a wide range of potential applications. The ability to accurately track discontinuous reference signals in a decentralized manner provides a natural fit for ancillary services, including primary and secondary frequency response. The efficacy of various control frequency control services that can be implemented on the back of this control mechanism is investigated in [26].

Moreover, the exact mapping of the aggregate TCL population onto a “leaky battery” representation [17] means that both response and recharging can be scheduled optimally according to aggregator objectives. Trovato *et al.* [27] have used this to optimally schedule primary and secondary response in the Great Britain system and calculate the associated benefits. The ability of some TCLs (e.g., commercial freezers) to significantly shift their energy consumption also opens up significant opportunities in energy arbitrage. An investigation of the optimal simultaneous allocation of multiple services (ancillary, arbitrage) for different TCL types is contained in [17].

The controller can therefore be applied to domestic, commercial, and industrial cooling and heating appliances with dynamics that are approximately described by (1), and that can be controlled in sufficient quantities to permit mean-field scheduling of the type (7). However, further work is required to extend the approach presented here to large refrigerators that have multiple compressors, as they do not strictly satisfy the on/off paradigm used here.

In future work, further analysis—both theoretical and simulation-based—will be done to investigate the robustness of the proposed control approach, including the effect of outdoor temperatures [25] and lockout constraints [13], [14]. This may result in enhancements of the controller that improve its performance in a wide range of scenarios. For example, 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. Finally, we highlight that there is a substantial research need to define common test cases and to use these to compare the performance of this and other aggregate TCL control approaches published in recent years.

REFERENCES

- [1] D. S. Callaway and I. A. Hiskens, “Achieving Controllability of Electric Loads,” *Proc. IEEE*, vol. 99, no. 1, pp. 184–199, Jan. 2011.
- [2] J. Vorwerk, U. Markovic, P. Aristidou, E. Vrettos, and G. Hug, “Modelling of variable-speed refrigeration for fast-frequency control in low-inertia systems,” Jun. 2020, *arXiv preprint, arXiv:2006.14410*. [Online]. Available: <https://arxiv.org/abs/2006.14410>
- [3] F. Schweppe, “Frequency adaptive, power-energy re-scheduler,” US Patent 4 317 049, 1979.
- [4] J. A. Short, D. G. Infield, and L. L. Freris, “Stabilization of grid frequency through dynamic demand control,” *IEEE Trans. Power Syst.*, vol. 22, no. 3, pp. 1284–1293, Aug. 2007.
- [5] D. Angeli and P.-A. Kountouriotis, “A stochastic approach to “Dynamic-Demand refrigerator control,” *IEEE Trans. Control Syst. Technol.*, vol. 20, no. 3, pp. 581–592, May 2012.
- [6] D. Hirst, “Responsive Load System,” Patent GB 2361118 A, 2000.
- [7] E. Webborn, “Natural heterogeneity prevents synchronization of fridges with deterministic frequency control,” *IEEE Access*, vol. 7, pp. 130 206–130 214, 2019.
- [8] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, “Aggregate flexibility of thermostatically controlled loads,” *IEEE Trans. Power Syst.*, vol. 30, no. 1, pp. 189–198, Jan. 2015.
- [9] C. Chong and A. Debs, “Statistical synthesis of power system functional load models,” in *Proc. 18th IEEE Conf. Decis. Control Including Symp. Adaptive Processes*, 1979, vol. 18, pp. 264–269.
- [10] S. Ihara and F. Schweppe, “Physically based modeling of cold load pickup,” *IEEE Trans. Power App. Syst.*, vol. PAS-100, no. 9, pp. 4142–4150, Sep. 1981.
- [11] J. L. Mathieu, S. Koch, and D. S. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *IEEE Trans. Power Syst.*, vol. 28, no. 1, pp. 430–440, Feb. 2013.
- [12] S. H. Tindemans, V. Trovato, and G. Strbac, “Decentralized control of thermostatic loads for flexible demand response,” *IEEE Trans. Control Syst. Technol.*, vol. 23, no. 5, pp. 1685–1700, Sep. 2015.
- [13] E. Vrettos, C. Ziras, and G. Andersson, “Fast and reliable primary frequency reserves from refrigerators with decentralized stochastic control,” *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 2924–2941, Jul. 2017.
- [14] L. C. Totu, R. Wisniewski, and J. Leth, “Demand response of a TCL population using switching-rate actuation,” *IEEE Trans. Control Syst. Technol.*, vol. 25, no. 5, pp. 1537–1551, Sep. 2017.
- [15] M. Chertkov and V. Chernyak, “Ensemble of thermostatically controlled loads: Statistical physics approach,” *Sci. Rep.*, vol. 7, no. 1, Dec. 2017, Art. no. 8673.
- [16] A. Basic and S. Meyn, “Distributed control of thermostatically controlled loads: Kullback–Leibler optimal control in continuous time,” in *Proc. IEEE 58th Conf. Decis. Control.*, Dec. 2019, pp. 7258–7265.
- [17] V. Trovato, S. H. Tindemans, and G. Strbac, “Leaky storage model for optimal multi-service allocation of thermostatic loads,” *IET Gener., Transmiss. Distrib.*, vol. 10, no. 3, pp. 585–593, Feb. 2016.
- [18] 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 *Proc. IEEE Int. Energy Conf.*, Leuven, 2016, pp. 1–6.
- [19] D. Ponce de Leon Barido, S. Suffian, J. Rosa, E. Brewer, and D. M. Kammen, “Enabling micro-level demand-side grid flexibility in resource constrained environments,” in *Proc. IEEE/ACM 2nd Int. Conf. Internet-of-Things Des. Implementation*, Pittsburgh, PA, USA, 2017, pp. 233–246.
- [20] S. Tindemans and G. Strbac, “Low-complexity control algorithm for decentralised demand response using thermostatic loads,” in *Proc. IEEE Int. Conf. Environ. Elect. Eng. Ind. Commercial Power Syst. Europe*, Jun. 2019, pp. 1–6.
- [21] S. H. Tindemans, V. Trovato, and G. Strbac, “Frequency control using thermal loads under the proposed ENTSO-E demand connection code,” in *Proc. IEEE Eindhoven PowerTech.*, Jun. 2015, pp. 1–6.

- [22] S. Tindemans, "Code release for Low-complexity algorithm for decentralized aggregate load control of thermostatic loads," 2020. [Online]. Available: <https://github.com/simontindemans/TCLcontrol/tree/v1.1>
- [23] S. Tindemans, "Compute capsule for Low-complexity algorithm for decentralized aggregate load control of thermostatic loads," 2020. [Online]. Available: <https://codeocean.com/capsule/6535638/tree/v1>
- [24] E. C. Kara, M. Berges, and G. Hug, "Impact of disturbances on modeling of thermostatically controlled loads for demand response," *IEEE Trans. Smart Grid*, vol. 6, no. 5, pp. 2560–2568, Sep. 2015.
- [25] K. Hreinsson, A. Scaglione, and M. Alizadeh, "An aggregate model of the flexible energy demand of thermostatically controlled loads with explicit outdoor temperature dependency," in *Proc. 53rd Hawaii Int. Conf. Syst. Sci.*, Jan. 2020, pp. 3075–3084.
- [26] V. Trovato, I. M. Sanz, B. Chaudhuri, and G. Strbac, "Advanced control of thermostatic loads for rapid frequency response in Great Britain," *IEEE Trans. Power Syst.*, vol. 32, no. 3, pp. 2106–2117, May 2017.
- [27] V. Trovato, F. Teng, and G. Strbac, "Role and benefits of flexible thermostatically controlled loads in future low-carbon systems," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 5067–5079, Sep. 2018.

Simon H. Tindemans (Member, IEEE) received the M.Sc. degree in physics from the University of Amsterdam, Amsterdam, The Netherlands, in 2004 and the Ph.D. degree in theoretical biophysics from Wageningen University, Wageningen, The Netherlands, in 2009.

From 2010 to 2017, he was a Research Associate, Marie-Curie Intra-European Fellow and Research Fellow with the Control and Power Group, Imperial College London, U.K. He is currently an Assistant Professor with the Department of Electrical

Sustainable Energy, Delft University of Technology, Delft, The Netherlands. His research focuses on statistical risk management for power systems. He has a specific research interest in the control of the aggregate power consumption of heterogeneous and unpredictable resources.

Goran Strbac (Member, IEEE) is a Professor of Energy Systems with Imperial College London, London, U.K. His current research interests include modeling and optimization of economic and security performance of energy system operation and investment, including analysis of future energy markets to support cost-effective transition to smart low carbon energy future. Prof. Strbac is currently Leading Author in IPCC WG 3, Member of the UK Smart System Forum, Member of the European Technology and Innovation Platform for Smart Networks for the Energy Transition, and Member of the Joint EU Programme in Energy Systems Integration of the European Energy Research Alliance.