Statistical mechanics of thermostatically controlled multizone buildings

Lucas Fuentes Valenzuela,¹ Lindell Williams,^{2,3} and Michael Chertkov³

¹Department of Electrical Engineering, Stanford University, 350 Jane Stanford Way Stanford, California 94305, USA

²Department of Physics, Cornell University, 109 Clark Hall, Ithaca, New York 14853, USA

³Program in Applied Mathematics & Department of Mathematics, University of Arizona, 617 N. Santa Rita,

Tucson, Arizona 85721, USA

(Received 27 August 2022; accepted 2 March 2023; published 29 March 2023)

We study the collective phenomena and constraints associated with the aggregation of individual cooling units from a statistical mechanics perspective. These units are modeled as thermostatically controlled loads (TCLs) and represent zones in a large commercial or residential building. Their energy input is centralized and controlled by a collective unit—the air handling unit (AHU)—delivering cool air to all TCLs, thereby coupling them together. Aiming to identify representative qualitative features of the AHU-to-TCL coupling, we build a simple but realistic model and analyze it in two distinct regimes: the constant supply temperature (CST) and the constant power input (CPI) regimes. In both cases, we center our analysis on the relaxation dynamics of individual TCL temperatures to a statistical steady state. We observe that while the dynamics are relatively fast in the CST regime, resulting in all TCLs evolving around the control set point, the CPI regime reveals the emergence of a bimodal probability distribution and two, possibly strongly separated, timescales. We observe that the two modes in the CPI regime are associated with all TCLs being in the same low or high airflow states, with an occasional collective transition between the modes akin to Kramer's phenomenon in statistical physics. To the best of our knowledge, this phenomenon has been overlooked in building energy systems despite its direct operational implications. It highlights a trade-off between occupational comfort—related to zonal temperature variations—and energy consumption.

DOI: 10.1103/PhysRevE.107.034140

I. INTRODUCTION

The essence of demand response (DR) lies in providing auxiliary services helping power system operators manage uncertainty. The latter can emerge both from variable generation, e.g., as wind and solar, or from electricity market volatility [1]. DR consists in leveraging flexible and inexpensive resources on the demand side of the power balance to ensure stability. The essential source of flexibility comes from many consumers of electricity tolerating consumption delays, provided that some constraints remain satisfied [2]. In addition to large and stable loads, aggregations of many small loads, e.g., residential appliances, can also be involved in DR services [3]. The heating and cooling system in residential buildings is one such load that possesses inherent flexibility due to thermal inertia, and thereby presents significant opportunities in a DR market. Specifically, the potential for impacting total load via set-point changes has been investigated numerically [4,5] and empirically [6–8], as well as its ramifications for occupant comfort [9].

A. Related work

Theoretical studies on the matter have centered around so-called thermostatically controlled loads (TCLs), denoting physical entities whose temperature oscillates within a range or around a target value; examples include rooms in buildings or refrigerators. Understanding the behavior of aggregations of individual entities and the underlying potential for DR has underpinned most of the interest in this discipline.

Initial studies focused on adapting existing and developing unique approaches to the statistics of TCLs [10,11], and proposed a methodology for the aggregation of individual loads [12]. Large aggregates, also called ensembles, were studied with the tools of statistical physics, such as Fokker-Planck equations [13,14], and of control and reinforcement learning, such as Markov decision processes (MDPs) [15,16]. The main operational philosophies in the literature include individual thermostat set-point control [17–19] or randomization and automatic feedback control at the individual TCL level, based on collective output [20,21].

While standard TCL models do represent some installations like independent AC units or refrigerators, they do not directly capture the intricate dynamics emerging from the coupling of those units via a district heating network [22] or the grid services they provide [17,23,24]. In the context of heating and cooling within a multizone building, individual zones are thermally regulated by a small number of air handling units (AHUs), see Fig. 1 for a simplified illustration. Each AHU is connected to some number of zones. Its role is to cool and dehumidify a mix of outside and recirculated air to a given temperature and relative humidity, and then to circulate this air-at a given temperature, the supply temperature-throughout the building. The airflow is distributed to each zone via variable air volume (VAV) boxes. See Ref. [25] for an accurate and detailed description of the operations of such systems. Standard TCL models fail to



FIG. 1. Simplified illustration of the setup considered. The AHU is represented by the fan, while individual units have their own temperatures. The exhaust system from each unit is not represented for simplicity, and air recirculation is neglected.

account for the specificities of such systems in two ways. First, they usually assume units to be independent. However, they are actually *coupled* through an AHU via the supply temperature. This interaction and its implication to building control has not been studied extensively in the TCL literature. On the contrary, most studies accounting for such coupling have centered on the development of algorithmic strategies [26-31]. While these perspectives are necessary for applications, in this paper we focus on the dynamical implications of those interactions. Second, the dynamics of cooling in these models is very different from the reality of residential and commercial buildings. Indeed, in standard TCL models, cycling is inherent in the dynamics. A unit is cooled at constant power until it crosses the lower limit of a deadband, and it is then left to reheat until it crosses its upper limit, at which point the cycle restarts. In AHU systems, such cycling is not enforced directly, and the cooling power changes following changes in the supply temperature [32]. We will see below that the absence of a prescribed deadband has far-reaching consequences on how fast the system can recover from a perturbation, e.g., caused by the units following a DR request.

B. Contribution

In this paper, we suggest models to represent TCLs coupled via an AHU, and study the thermalization dynamics of the system subject to two-level control—at the collective (AHU) and at the individual (TCL) level. Specifically, we consider two principally different models representing distinct scenarios.

1. Constant supply temperature (CST). The supply air temperature cooling each TCL is held constant over time, thereby eliminating direct coupling between individual units. While simplistic, this model enables the understanding of the fundamental dynamics of TCL units.

2. *Constant power input (CPI).* TCLs are coupled via the supply temperature as is the case in large residential and commercial buildings. The latter is adjusted to ensure constant power input (CPI).

In both cases, we focus on analyzing the details of the stabilization of the ensemble to a statistical steady state, established as a result of the balance between exogenous fluctuations representing TCL-specific perturbations and thermal relaxation. We combine analytical and numerical analyses to derive insights into the dynamics of this system. In the CST model, we find that thermalization is relatively fast and results in TCLs stabilizing around the prescribed temperature. However, in the CPI scenario, we observe the appearance of more complex dynamics. When the level of stochastic fluctuations is small, all units stabilize in either of the two states associated with low or high airflow. Transitions between both modes, and thus relaxation to a statistical steady state, happens at a timescale which is much longer than the natural timescale of the problem, associated with a metastable equilibrium around one of the modes. Increasing the level of fluctuations results in a phase transition to an entropic state distributed around the control set point. We argue that this complex behavior of the CPI model is akin (but not exactly equivalent) to the phenomenon of the thermally activated barrier crossing in statistical mechanics, often referred to as the reaction rate, or Kramers' theory [33–36].

These observations are directly relevant for the operations of those common coupled systems. Under low uncertainty and perturbations, one can argue that operating the system at fixed supply temperature (CST) will result in a predictable outcome matching comfort requirements, with minor fluctuations in the total power required. However, in the case of higher fluctuations, e.g., associated with other uncertainties or driving forces such as sun irradiation, the nontrivial dynamics imply a need to balance comfort and energy requirements, e.g., for use in DR.

C. Outline

The contents are presented as follows. In Sec. II, the problem is formulated along with the two scenarios of interest. A statistical analysis of the CST scenario is conducted in Sec. III. Analyses under CPI are conducted in Sec. IV. We first consider the noiseless limit for one TCL, and extend its insights to the presence of noise. The general case with an arbitrary number of TCLs is studied in Sec. IV C. We conclude and discuss the path forward in Sec. V.

II. PROBLEM FORMULATION

The system we consider consists of two levels. At the lower level, TCLs evolve according to their individual dynamics. At the higher level, the AHU manages the energetic input to all TCLs, thereby coupling them to one another (cf. Fig. 1). In this section, we lay out the models governing both levels.

A. TCL dynamics

We model the dynamics of an individual TCL by the following stochastic differential equation (SDE):

$$c dT_i = -f(T_i(t); s(T_i(t)); T_s) dt + \sqrt{2D} dW_i(t),$$
 (1)

$$f(T_i(t); s(T_i(t)); T_s) = \frac{T_i(t) - T_o}{r} + \bar{\mu} s(T_i(t)) c_p(T_i(t) - T_s), \quad (2)$$

$$s(T_i(t)) = \begin{cases} s_-, & T_i(t) \leq \bar{T} \\ s_+, & T_i(t) > \bar{T}, \end{cases}$$
(3)

where t is time and $T_i(t)$ is the temperature of the *i*-th TCL unit $i \in 1, ..., N$. We assume that all units are identical and characterized by the following physical parameters: r [K/kW]is the coefficient of thermal resistance to the outside air; c_p [kJ/kgK] is the specific heat capacity of air; c [kJ/K] is the capacitance coefficient representing the total air-mass associated with an individual TCL; $W_i(t)$ is the Wiener process; and D [kJ²] represents the amplitude of exogenous, stochastic, zero mean thermal (white) noise affecting the TCLs (associated, e.g., with some uncertainty around in-zone traffic and operations). $\bar{\mu}$ denotes the maximum airflow that can get into the zone, while $s(\cdot) \in [0, 1]$ indicates how much air is actually flowing into the room at time t. In Eq. (2), describing the temperature dynamics of a TCL unit, $f(T_i(t); s(T_i(t)); T_s)$ denotes the thermal force associated with two principally different terms. The first contribution to the thermal force describes deterministic relaxation of the zone temperature, $T_i(t)$, with the rate r to the ambient temperature T_o , which is taken as constant in this paper. The second contribution to the thermal force represents the injection of cool air delivered by VAV units at the mass flow rate $\bar{\mu}s(T_i(t))$ and at a temperature equal to the supply temperature T_s into the *i*-th zone. In the present system, we assume the VAVs are able to modify only the airflow entering each zone. While in general they are also able to reheat the incoming air, we assume this is not applicable, as is the case in warm regions where cooling is mostly needed. The dynamics of the airflow at each VAV is a simplification of the dual-maximum control logic [37]. We neglect the heating regime and assume instantaneous switching (s(t) in Eq. (3))as the set point is crossed instead of gradual changes in the airflow.

The TCL Eqs. (1)–(3), as well as those describing the AHU dynamics discussed in the next subsection, are reduced in the sense that all the transients within the TCL units as well as within the AHU are ignored. This thermodynamic modeling is justified because the transients processes-which would require a detailed modeling of the energy transfersare largely complete within minutes, while we are interested in describing the dynamics of temperature within the building over longer timescales (tens of minutes to hours). Consistently with the consideration of the time separation as well as of universality, we model uncertainty due to thermal forces exerted on an individual TCL as white noise. This simplification allows us to write Fokker-Planck equations and then to derive analytical solutions for some probability distribution functions of interest. We do not expect qualitatively different conclusions in the case of a more realistic colored noise, even though the present derivations do not generalize.

B. AHU dynamics

The model in Eqs. (1)–(3) is incomplete until one provides a closure relationship for T_s , controlling the second contribution to the thermal force in Eq. (2). As detailed in the remainder of this section, we consider two scenarios which will be different in terms of the relation between T_s and the vector of TCL temperatures, $T = [T_1, ..., T_N]$. In this paper, we consider simplified models of the AHU energy consumption. Indeed, we assume that the units lie in areas where the energy attributed to cooling dominates that associated to dehumidification. This means that the energy consumption of the AHU can be modeled by the enthalpy change of the air as it goes through the AHU [38–41]. For the sake of simplicity, we also neglect air recirculation within the building and assume that only the outside air is used as input into the AHU.

1. Constant supply temperature. The first scenario considered, indexed by a, consists of fixing the temperature T_s of the air delivered to the TCLs to a prescribed constant value, $T_{s;a}$. The total power, i.e., energy injected into the ensemble per unit time, then writes

$$P_a(t) = \mu(t)(T_o - T_{s;a}), \quad \mu(t) = \bar{\mu} \sum_i s(T_i(t)), \quad (4)$$

where $\mu(t)$ is the aggregate airflow within the AHU. We assume that all the airflows in the system are balanced at any moment in time. We note that under this CST scenario, where $T_s = \text{const}, s(T_i(t))$ changes with time due to the evolution of T_i , thereby driving changes in $\mu(t)$. The total power consumed by the ensemble, P(t), therefore inherits temporal dynamics from the aggregation of individual TCL behaviors.

2. Constant power input. In this second scenario, indexed by b, we assume that the system operator keeps the power input, P, constant at all times, i.e., $P = P_b = \text{const.}$ This is achieved by adjusting the supply temperature, $T_s(t)$, according to the following modification of Eq. (4):

$$T_{s;b}(\boldsymbol{T}(t)) = T_o - \frac{P_b}{\bar{\mu} \sum_i s(T_i(t))}.$$
(5)

We see that the supply temperature is continually adjusted based on the local temperature dynamics.

III. STATISTICAL ANALYSIS OF THE CST SCENARIO

The supply temperature T_s is assumed constant, $T_s = T_{s;a}$. The system of SDEs (1)–(3) is closed, thus translating into the so-called Kolmogorov-Fokker-Planck (KFP) [42,43] partial differential equation for the joint probability distribution of the vector of temperatures T within the TCL ensemble:

$$\partial_t \mathscr{P}_a(\boldsymbol{T}|t) = \frac{1}{c} \sum_i \; \partial_{T_i}(f(T_i; s(T_i); T_{s;a}) + D\partial_{T_i}) \mathscr{P}_a(\boldsymbol{T}|t).$$
(6)

We observe that the KFP Eq. (6) can be represented in the potential form

$$\partial_t \mathscr{P}_a(\boldsymbol{T}|t) = -\sum_i \partial_{T_i} J_i,\tag{7}$$

$$J_{i} = -\frac{1}{c} (\partial_{T_{i}} U(\mathbf{T}|T_{s;a}) + D\partial_{T_{i}}) \mathscr{P}_{a}(\mathbf{T}|t), \qquad (8)$$

$$U(T|T_{s;a}) = \sum_{i=1}^{N} U_1(T_i|T_{s;a}),$$
(9)

$$U_{1}(T_{i}|T_{s;a}) = \frac{(T_{i} - T_{o})^{2}}{2r} + \frac{c_{p}\bar{\mu}s(T_{i})}{2} ((T_{i} - T_{s;a})^{2} - (\bar{T} - T_{s;a})^{2}), \quad (10)$$

where J_i is the probability current along T_i , and $U(T|T_{s;a})$ and $U_1(T_i|T_{s;a})$ are the aggregated and individual TCL thermal potentials, respectively. As the KFP Eq. (6) describes an initial-value problem, it should be equipped with an initial



FIG. 2. Temperature distributions for an individual TCL under the CST scenario after initial transients for different values of the supply temperature T_s . c = 20.5 kJ/K, r = 2K/kW, $D = 25\text{kJ}^2$, $\overline{T} = 25 \text{ C}$, $(s_-, s_+) = (0.2, 1)$, $\mu = 1 \text{ kg/s}$. Filled: Experimental results. Dotted lines: Theoretical distributions obtained according to Eq. (11).

condition, e.g., that all TCLs are at the same temperature T_0 at t = 0. Then, $\mathcal{P}_a(T|0) = \prod_i \delta(T_i - T_0)$, where $\delta(\cdot)$ is the Dirac δ function. More generally, if $\mathcal{P}_a(\mathbf{T}|0)$ is factorized into a product of marginal distributions of individual T_i , the solution to the KFP equations at all t is also factorized into the product of the corresponding marginal probability distributions, i.e., $\mathcal{P}_a(\mathbf{T}|t) = \prod_i \mathcal{P}_{a;1}(T_i|t)$.

A. Stationary distribution

Due to the factorized structure of the differential operator on the right-hand side of Eq. (6), the steady-state solution to Eq. (6), $\mathcal{P}_{a;st}(T)$, i.e., one acquired at $t \to \infty$, is also factorized into the product of the respective marginals, $\mathcal{P}_{a;st}(T) =$ $\prod_i \mathcal{P}_{a;1;st}(T_i|T_{s;a})$. Moreover, $\mathcal{P}_{a;1;st}(T_i|T_{s;a})$ satisfies a second order ordinary differential equation (ODE) which can be solved for any values of $T_{s;a}$, considered as a parameter:

$$\mathscr{P}_{a;1;st}(T_i|T_{s;a}) = \frac{1}{Z} \exp\left(-\frac{cU_1(T_i|T_{s;a})}{D}\right), \qquad (11)$$

$$Z = \int_{-\infty}^{\infty} dT \exp\left(-\frac{cU_1(T|T_{s;a})}{D}\right).$$
 (12)

The single-TCL thermal potential, $U_1(T_i|T_{s;a})$, defined in Eq. (9), is continuous in T_i . It is, however, not smooth as the derivative jumps at $T_i = \overline{T}$, and attains a single minimum at

$$T_{a;\min}(T_{s;a}) = \begin{cases} \frac{T_o + T_{s;a}\mu^{(-)}cr}{1+\mu^{(-)}cr}, & T_{s;a} \leqslant \beta^{(-)} \\ \bar{T}, & \beta^{(-)} \leqslant T_{s;a} \leqslant \beta^{(+)} \\ \frac{T_o + T_{s;a}\mu^{(+)}cr}{1+\mu^{(+)}cr}, & \beta^{(+)} \leqslant T_{s;a}, \end{cases}$$
(13)

where $\beta^{(\pm)} = \overline{T} - (T_o - \overline{T})/(\mu^{(\pm)}cr)$. The most probable value of the TCL temperature at steady state does not depend on the amplitude of the thermal noise, *D*.

We report in Fig. 2 the dependence of the steady-state individual TCL probability distribution $\mathcal{P}_{a;1;st}(T_i|T_{s;a})$ on T_i at different values of $T_{s;a}$. Throughout this paper, unless noted



FIG. 3. Location of the two fixed points for N = 2 in the CPI scenario. The identity is represented by the gray diagonal line. Green line, dashed: Expected position of the two fixed points according to Eq. (14). Blue dots: Final temperatures of both TCLs attained after 100s under multiple initializations.

otherwise, parameters for numerical results are $T_o = 30$ °C, $s_- = 0.2$, $s_+ = 1$, $\bar{\mu} = 1$ kg/s, r = 2 K/kW, c = 15 kJ/K, $\bar{T} = 25$ °C, $T_s = 10$ °C. Note that there is a range of supply temperature values where the most probable value is achieved at $T_i = \bar{T}$, e.g., for $T_s = 20$ °C in Fig. 2. This maximum appears because fluctuations drive the temperature across \bar{T} , thereby implying repetitive switching of the airflow $s(T_i)$, and on average stabilizing the temperature around the set point.

B. Relaxation to the steady distribution

Even in the case when the initial probability distribution cannot be factorized into a product of independent terms, independence will appear dynamically at sufficiently large $t \gg$ τ_a , where τ_a is the so-called mixing time [44]. One can extract τ_a and, specifically, its dependence on the supply temperature $T_{s;a}$ from the spectral analysis of the the *i*-th component of the differential operator entering the KFP Eq. (6). The resulting dependence does not hold any specific and unexpected features and is therefore not reported here.

IV. STATISTICAL ANALYSIS OF THE CPI SCENARIO

The most significant difference between the CST and the CPI scenarios lies in the stochastic coupling between individual TCLs. According to Eq. (5), the supply temperature T_s directly depends on all TCL temperatures. In this section, we discuss the significance of this modification.

Informally, we expect that this case is special because the stochastic dynamics of individual TCLs are no longer decoupled. The general theory of KFP equations suggests [42,43] that in the CPI scenario with N > 1, (a) the thermal force can no longer be represented as a gradient of a potential; (b) detailed balance is broken; and (c) the steady distribution is no longer a Gibbs distribution and cannot be factorized into a product of components, each representing an individual TCL.

In this section, we provide a quantitative analysis of the aforementioned expected qualitative behavior of the system. In Sec. IV A, we analyze the system in the noiseless limit, i.e., D = 0, and show that the system is strongly sensitive to initial conditions via the identification of two different fixed points governing the long-time behavior of the system. In Sec. IV B, we reintroduce noise in the context of a single TCL as it enables some analytical treatment. We leverage the latter in



FIG. 4. Dynamics and steady state distribution of one TCL under the CPI scenario. (a) Temperature distribution of 1 TCL under the CPI scenario after initial transients, for different values of the noise amplitude *D*. Filled: Experimental results. Dashed line: Theoretical distributions obtained according to Eq. (11) and (b) Time series of 1 TCL under the CPI scenario, with $D_i = 15 \text{ kJ}^2$. Red lines, dashed: Locations of the theoretical fixed points T_{\pm} [cf. Eq. (14)]. Filled background: A change in background color indicates a change in the value of $T_{\min}(T_s)$.

Sec. IV C where the general case of N > 1 TCLs with noise is considered and where we discover that several features from the N = 1 case surprisingly hold.

A. Noiseless limit

To extract generic statements about the behavior of this coupled system, we analyze the fixed points of Eq. (1) with T_s substituted by $T_{s;b}(T(t))$ from Eq. (5) in the noiseless case, i.e., in the deterministic regime. As TCLs are indistinguishable, a fixed point solution is determined by the number of TCLs in the high-airflow regime, $N_+(t) = \sum_i \theta(T_i(t) > \overline{T})$, where $\theta(\cdot)$ is the indicator function. We observe that there are potentially N + 1 fixed point solutions to Eq. (1) complemented with the power balance in the CPI scenario, i.e., $N_+(t) \in \{0, \ldots, N\}$. According to Eq. (5), $T_{s;b}$ is directly determined by $N_+(t)$: $T_{s;b}(N_+(t)) = T_o - P_b/\overline{\mu}(N_+s_+ + (N - C_{s;b}))$



FIG. 5. Temperature distribution and dynamics in the general case of N > 1, D > 0 CPI case. (a) Empirical temperature distributions of TCLs under the CPI scenario with N = 5, after initial transients and for different values of the noise amplitude D. (b) Time series of two TCLs under the CPI scenario, with N = 2. Red lines, dashed: Locations of the theoretical fixed points T_{\pm} . Filled background: A change in background color indicates a switch in the value of $T_{\min}(T_s)$.

 $N_+)s_-$). For each of the N + 1 possible states, the N TCLs are split into two groups: a group of N_+ units in the high-flow regime and another $N - N_+$ in the low-flow regime. In addition, in the deterministic case, all TCLs in a given group are at the same temperature at the steady state, given by

$$T_{\pm}(N_{+}(t)) = \frac{T_{o} + T_{s;b}(N_{+}(t))\mu_{\pm}cr}{1 + \mu_{\pm}cr}.$$
 (14)

An immediate consequence of this equation is that $T_+(N_+(t)) < T_-(N_+(t))$, providing that $\mu_+ > \mu_-$. However, this is in contradiction with the requirement that $T_+(N_+(t)) > \overline{T} > T_-(N_+(t))$. In other words, there can be no stable fixed points, i.e., states realized dynamically at $t \to \infty$ in the noiseless regime, with a nonzero number of TCLs in both states (high flow and low flow). Therefore, two options are left for stable fixed points: $N_+(t) = 0$ or $N_+(t) = N$. Which of the two stable fixed points is attained at long



FIG. 6. Statistical analysis and observation of two different timescales in the system. (a) Autocorrelation of $T(t) - T_{\min}(t)$. The exponential decay is indicative of a short associated timescale. (b) Autocorrelation of T_{\min} , defined in (13). The decay is much slower than for Fig. 6(a), and the characteristic timescale increases with the number of TCLs in the system.

times will depend on the value of \overline{T} . Indeed, the $N_+(t) = 0$ state is valid (self-consistent) if $T_-(N_+(t) = 0) < \overline{T}$ and the $N_+(t) = N$ state is valid if $T_+(N_+(t)) > \overline{T}$. In addition, $T_-(N_+(t) = 0) < T_+(N_+(t) = N)$. Therefore, the steady state that will be attained will consist of (i) $N_+(t) = N$ if $\overline{T} < T_-(N_+(t) = 0) < T_+(N_+(t) = N)$, (ii) $N_+(t) = 0$, if $T_-(N_+(t) = 0) < T_+(N_+(t) = N) < \overline{T}$, or (iii) either of the two if $T_-(N_+(t) = 0) < \overline{T} < T_+(N_+(t) = N)$, in which case the final state is directly dictated by the initial conditions T(0). Figure 3 illustrates the dependence of the realized fixed points on the values of \overline{T} for different initial conditions in the case of N = 2 TCLs. T_{\pm} do indeed act as stable fixed points in the noiseless limit. Depending on \overline{T} , either one or both of them are accessible to the system.

B. Case of N = 1, D > 0

As analysis in the general case with N > 1, D > 0 is difficult, we focus here on the stochastic case for one TCL. Equations (10) and (11) still hold in this context, with the difference that the supply temperature is dictated by Eq. (5). A similar analysis as that conducted in Sec. III is hence possible, and results are reported in Fig. 4.

In contrast to the single-peak distributions observed in the CST scenario (cf. Fig. 2), one observes the emergence of two maxima corresponding to the two fixed points identified in Sec. IV A. The position of these stable temperatures is directly determined by the physical parameters of the system, while their relative importance can be tuned by the set point \overline{T} . We attribute the mismatch between analytical and empirical



FIG. 7. Scaling of the switching time with the number N of TCLs in the CPI scenario. Dots indicate the average switching time. Error bars indicate the 75th and 25th percentile of the distributions.

distributions in Fig. 4(a) to a finite computational budget and step size. The stochastic nature of the system actually enables switching between both modes. As illustrated in Fig. 4(b), a single TCL driven by noise successively switches between the two equilibrium points, thereby implying an adaptation of the supply air temperature T_s , not shown here.

As the potential $U_1(T)$ has two distinct minima in some regimes, Kramers theory [33–36] suggests the emergence of two distinct temporal scales. One short timescale τ_s should be associated with equilibrium fluctuations around either of the two minima, while the transition between either minimum as seen in Fig. 4(b) is characterized by a longer one τ_l . These are related via $\tau_l \sim \tau_s \exp(\Delta/D)$, where Δ is the potential barrier at $T_1 = \overline{T}$. Strong separation of these timescales, i.e., $\tau_l \gg \tau_s$, emerges when $D \ll \Delta$.

C. Numerical experiments for N > 1, D > 0

Even if, strictly speaking, insights from Kramers theory only apply to the N = 1 case discussed above, most qualitative features seem to translate to the general N > 1 case.

a. Bimodal dynamics and emergence of two timescales. Up to moderate values of the noise amplitude D a similar bimodal distribution of accessed temperatures is observed (cf. Fig. 5(a), with $D = 10 \text{ kJ}^2$). Oscillations of temperatures around two central values are clearly visible in Fig. 5(b). Note that both TCLs exhibit coordinated switching from one fixed point to the other, resulting from the AHU coupling. Both observations are qualitatively analogous to the N = 1case, cf. Fig. (4). Specifically, one can also expect switching and the natural appearance of two characteristic timescales in the system. As shown in Fig. 6(a), the autocorrelation $\rho(T(t) - T_{\min}(t))$ of a given TCL temperature decays exponentially fast. However, the decay of $\rho(T_{\min}(t))$, associated with the N TCLs switching from one stable point to the next, is significantly slower in Fig. 6(b). This suggests that memory is indeed present in the system and that switching is a rarer event with increasing values of N. Figure 7 provides another illustration of this strong N dependence. The switching time over trials increases with N, confirming the intuition provided by Fig. 6(b). While the dependence of the switching time with N seems subexponential, those results were obtained with a finite computational budget and might not reflect the true scaling. Further investigation is needed, via alternative theoretical and computational methods, to resolve it.

b. Emergence of an entropic state. Large noise amplitudes will prevent the system from settling in either of the stable points. Indeed, the two modes of the temperature distribution disappear as D increases [see Fig. 5(a)]. Instead, they are slowly replaced by an *entropic* peak at the location of the set point \overline{T} . From the perspective of a single TCL, this entropic state represents a balance between stochastic and deterministic forces resulting in the TCL switching back and forth between low- and high-flow regimes, and hence meandering around \overline{T} . When the amplitude of random fluctuations is high, the effect of the potential barrier disappears and the resulting distribution is akin in nature to the CST case in some regimes as shown in Fig. 2.

V. CONCLUSIONS AND PATH FORWARD

In this paper, we formulate a model describing the stochastic dynamics of coupled TCLs. It differs from standard TCL models in that it more accurately replicates cooling via a centralized AHU. The main consequence of this modification is the resulting coupling between individual units, which is usually disregarded in other models.

We analyzed the aggregate dynamics in two different scenarios. First, we considered the uncoupled dynamics of each unit by fixing the supply air temperature (CST regime). While the CST regime is highly stable and predictable, fluctuations of the aggregated power consumption of the building associated with airflow adjustments in each unit might be undesirable. In the second CPI scenario, the power input is set constant, leading to significant coupling between individual TCLs and revealing interesting collective phenomena, with ramifications for use of the ensemble in DR. The system's response to a perturbation may result in the coexistence of two quasi-steady modes associated with stable fixed points of the dynamics in the low-noise regime, which transitions to an *entropic* state when the level of the noise increases. A dynamic consequence of this complex behavior is the emergence of long transients associated with the transitions between the two modes in the low-noise regime.

This model and associated results set the stage for several future challenges:

(1) Investigating the effects associated with inhomogeneity (disorder) in physical parameters, thereby describing populations of nonidentical units, would provide a more accurate picture of real-life systems.

(2) Developing methods leveraging both unit-level and system-level control to strike a balance between individual comfort and collective benefits would further our understanding of the inherent flexibility in such systems and potentially result in innovative solutions for AHU control—e.g., to remove the anomalously long relaxation times observed in the low-noise CPI regime. Ideas in recent papers [20,21] could be translated to the present case. Specifically, understanding the nonequilibrium nature of the TCL ensemble in the CPI regime and its practical significance for operations and control will be interesting. While similar to typical Kramers phenomena, the dynamics in the CPI regime are not at statistical equilibrium. The detailed balance is broken and fluxes emerge. The implications of such complex dynamics on controllability remain to be explored.

(3) Confronting this model with real zone-level time series and experiments [45,46] will provide crucial insight for its validity and highlight where improvements, modifications, or relaxing assumptions are needed.

(4) Analyzing the integration of such systems into largerscale DR schemes, e.g., in the spirit of Refs. [47,48], and involving both multiple buildings as well as other appliances, will be critical in the deployment of such solutions.

ACKNOWLEDGMENTS

L.F.V. acknowledges Jacques de Chalendar for inspiring and helpful discussions, as well as for proofreading the paper. Research Experience for Undegraduates work of L.W. at University of Arizona during the summer of 2021 was funded by NSF No. 1937229: RTG: Applied Mathematics and Statistics for Data-Driven Discovery.

- P. Palensky and D. Dietrich, Demand side management: Demand response, intelligent energy systems, and smart loads, IEEE Trans. Ind. Inf. 7, 381 (2011).
- [2] N. Oconnell, P. Pinson, H. Madsen, and M. Omalley, Benefits and challenges of electrical demand response: A critical review, Renewable Sustainable Energy Rev. 39, 686 (2014).
- [3] R. D'hulst, W. Labeeuw, B. Beusen, S. Claessens, G. Deconinck, and K. Vanthournout, Demand response flexibility and flexibility potential of residential smart appliances: Experiences from large pilot test in Belgium, Appl. Energy 155, 79 (2015).
- [4] R. Yin, E. C. Kara, Y. Li, N. DeForest, K. Wang, T. Yong, and M. Stadler, Quantifying flexibility of commercial and residential loads for demand response using setpoint changes, Appl. Energy 177, 149 (2016).

- [5] C. Finck, R. Li, R. Kramer, and W. Zeiler, Quantifying demand flexibility of power-to-heat and thermal energy storage in the control of building heating systems, Appl. Energy 209, 409 (2018).
- [6] J. A. de Chalendar, C. McMahon, L. F. Valenzuela, P. W. Glynn, and S. M. Benson, Unlocking demand response in commercial buildings: Empirical response of commercial buildings to daily cooling set point adjustments, Energy Buildings 278, 112599 (2023).
- [7] P. Xu, P. Haves, and M. A. Piette, Demand shifting with thermal mass in large commercial buildings: Field tests, simulation and audits,, 35–43 (2010).
- [8] E. C. Kara, M. D. Tabone, J. S. MacDonald, D. S. Callaway, and S. Kiliccote, Quantifying flexibility of residential thermostatically controlled loads for demand response: A data-driven

approach, BuildSys 2014 - Proc. 1st ACM Conf. Embedded Syst. Energy-Efficient Buildings 140 (2014).

- [9] S. Aghniaey and T. M. Lawrence, The impact of increased cooling setpoint temperature during demand response events on occupant thermal comfort in commercial buildings: A review, Energy Build. 173, 19 (2018).
- [10] C. Y. Chong and A. S. Debs, Statistical synthesis of power system function load models, Proc. IEEE Conf. Decision Control 1, 264 (1979).
- [11] S. Ihara and F. C. Schweppe, Physically based modeling of cold load pickup, IEEE Trans. Power Appar. Syst. PAS-100, 4142 (1981).
- [12] R. Malhamé and C. Y. Chong, Statistical synthesis of physically based load models with applications to cold load pickup, IEEE Trans. Power Appar. Syst. PAS-103, 1621 (1984).
- [13] R. Malhamé and C. Y. Chong, Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system, IEEE Trans. Automat. Contr. 30, 854 (1985).
- [14] R. Malhamé and C.-Y. Chongt, On the statistical properties of a cyclic diffusion process arising in the modeling of thermostatcontrolled electric power system loads, SIAM J. Appl. Math. 48, 465 (1988).
- [15] M. Chertkov and V. Chernyak, Ensemble of thermostatically Controlled Loads: Statistical physics approach, Sci. Rep. 7, 8673 (2017).
- [16] M. Chertkov, D. Deka, and Y. Dvorkin, Optimal ensemble control of loads in distribution grids with network constraints, in 2018 Power Systems Computation Conference (PSCC) (2018), pp. 1–7.
- [17] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, Energy Convers. Manage. 50, 1389 (2009).
- [18] D. S. Callaway and I. A. Hiskens, Achieving controllability of electric loads, Proc. IEEE 99, 184 (2011).
- [19] S. Bashash and H. K. Fathy, Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads, Proc. Am. Control Conf., 4546 (2011).
- [20] D. Métivier, I. Luchnikov, and M. Chertkov, Power of ensemble diversity and randomization for energy aggregation, Sci. Rep. 9, 5910 (2019).
- [21] D. Métivier and M. Chertkov, Mean-field control for efficient mixing of energy loads, Phys. Rev. E 101, 022115 (2020).
- [22] B. J. Claessens, D. Vanhoudt, J. Desmedt, and F. Ruelens, Model-free control of thermostatically controlled loads connected to a district heating network, Energy Build. 159, 1 (2018).
- [23] E. Webborn and R. S. Mackay, A stability analysis of thermostatically controlled loads for power system frequency control, Complexity 2017, 1 (2017).
- [24] S. Bashash and H. K. Fathy, Modeling and control of aggregate air conditioning loads for robust renewable power management, IEEE Trans. Control Syst. Technol. 21, 1318 (2013).
- [25] Y. Ma, A. Kelman, A. Daly, and F. Borrelli, Predictive control for energy efficient buildings with thermal storage, IEEE Control Syst. 32, 44 (2012).
- [26] T. I. Salsbury, A temperature controller for VAV air-handling units based on simplified physical models, HVAC R Res. 4, 265 (1998).

- [27] S. Wang and X. Jin, Model-based optimal control of VAV airconditioning system using genetic algorithm, Build. Environ. 35, 471 (2000).
- [28] S. Ma, Y. Zou, and S. Li, Coordinated control for air handling unit and variable air volume boxes in multi-zone HVAC system, J. Process Control 107, 17 (2021).
- [29] W. Li, S. Wang, and C. Koo, A real-time optimal control strategy for multi-zone VAV air-conditioning systems adopting a multi-agent based distributed optimization method, Appl. Energy 287, 116605 (2021).
- [30] J. M. Lee, S. H. Hong, B. Mo Seo, and K. H. Lee, Application of artificial neural networks for optimized AHU discharge air temperature set-point and minimized cooling energy in VAV system, Appl. Thermal Engineer. 153, 726 (2019).
- [31] A. Kelman and F. Borrelli, Bilinear model predictive control of a HVAC system using sequential quadratic programming, IFAC Proc. Vol. 44, 9869 (2011).
- [32] J. S. MacDonald, E. Vrettos, and D. S. Callaway, A critical exploration of the efficiency impacts of demand response from HVAC in commercial buildings, Proc. IEEE 108, 1623 (2020).
- [33] H. Kramers, Brownian motion in a field of force and the diffusion model of chemical reactions, Physica 7, 284 (1940).
- [34] P. Hänggi, P. Talkner, and M. Borkovec, Reaction-rate theory: Fifty years after Kramers, Rev. Mod. Phys. 62, 251 (1990).
- [35] E. Pollak, Variational transition state theory for activated rate processes, J. Chem. Phys. 93, 1116 (1990).
- [36] V. Mel'nikov, The Kramers problem: Fifty years of development, Phys. Rep. 209, 1 (1991).
- [37] S. T. Taylor, J. Stein, G. Paliaga, and H. Cheng, Dual maximum VAV box control logic, ASHRAE J. 54, 16 (2012).
- [38] R. Sterling, G. Provan, J. Febres, D. O'Sullivan, P. Struss, and M. M. Keane, Model-based fault detection and diagnosis of air handling units: A comparison of methodologies, in *Energy Procedia* (2014), Vol. 62.
- [39] A. Razban, A. Khatib, D. Goodman, and J. Chen, Modelling of air handling unit subsystem in a commercial building, Therm. Sci. Eng. Prog. 11, 231 (2019).
- [40] I. Zajic, T. Larkowski, M. Sumislawska, K. J. Burnham, and D. Hill, Modelling of an air handling unit: A Hammerstein-bilinear model identification approach, in 2011 21st International Conference on Systems Engineering, Las Vegas, NV, USA (IEEE, 2011), pp. 59–63.
- [41] A. Afram and F. Janabi-Sharifi, Review of modeling methods for HVAC systems, Appl. Therm. Eng. 67, 507 (2014).
- [42] C. W. Gardiner, Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences (Springer-Verlag, Berlin, 2004), Vol. 13.
- [43] N. G. Van Kampen, in North-Holland Personal Library, Stochastic Processes in Physics and Chemistry, 3rd ed. (Elsevier, 2007).
- [44] τ_a is defined as a time when $\mathscr{P}_{a;1}(T|t; T_{s;a})$ gets sufficiently close, wrt norm L_2 or Kullback-Leibler divergence, to its steady asymptotics, i.e., $\|\mathscr{P}_{a;1;st}(T|T_{s;a}) \mathscr{P}_{a;1}(T|t; T_{s;a})\| \propto \exp(-t/\tau_a) \ll 1.$
- [45] A. Keskar, S. Lei, T. Webb, S. Nagy, H. Lee, I. A. Hiskens, J. L. Mathieu, and J. X. Johnson, Stay cool and be

flexible: Energy-efficient grid services using commercial buildings HVAC systems, in *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings* (2020).

- [46] P. Li, D. Vrabie, D. Li, S. C. Bengea, S. Mijanovic, and Z. D. O'Neill, Simulation and experimental demonstration of model predictive control in a building HVAC system, Sci. Technol. Built Environ. 21, 721 (2015).
- [47] A. Hassan, R. Mieth, M. Chertkov, D. Deka, and Y. Dvorkin, Optimal load ensemble control in chance-constrained optimal power flow, IEEE Trans. Smart Grid 10, 5186 (2019).
- [48] A. Hassan, S. Acharya, M. Chertkov, D. Deka, and Y. Dvorkin, A hierarchical approach to multienergy demand response: From electricity to multienergy applications, Proc. IEEE 108, 1457 (2020).