



HAL
open science

Aggregate capacity for TCLs providing virtual energy storage with cycling constraints

Austin Robert Coffman, Ana Bušić, Prabir Barooah

► **To cite this version:**

Austin Robert Coffman, Ana Bušić, Prabir Barooah. Aggregate capacity for TCLs providing virtual energy storage with cycling constraints. CDC 2019 - 58th IEEE Conference on Decision and Control, Dec 2019, Nice, France. hal-02425977

HAL Id: hal-02425977

<https://hal.science/hal-02425977>

Submitted on 31 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Aggregate capacity for TCLs providing virtual energy storage with cycling constraints

Austin R. Coffman^{*,‡} Ana Bušić[†] Prabir Barooah^{*}

^{*}University of Florida {bubbaroney, pbarooah}@ufl.edu [†]INRIA ana.busic@inria.fr

Abstract—The coordination of thermostatically controlled loads (TCLs) is challenging due to the need to meet individual loads quality of service (QoS), such as indoor temperature constraints. Since these loads are usually on/off type, cycling rate is one of their QoS metrics; frequent cycling between on and off states is detrimental to them. While significant prior work has been done on the coordination of air conditioning TCLs, the question of cycling QoS has not been investigated in a principled manner. In this work we propose a method to characterize aggregate capacity of a collection of air conditioning TCLs that respects the loads cycling rate constraints (maximum number of cycles in a given time period). The development is done within the framework of randomized local control in which a load makes on/off decisions probabilistically. This characterization allows us to propose a reference planning problem to generate feasible reference trajectories for the ensemble that respect cycling constraints. The reference planning problem manifests itself in the form a Nonlinear Programming problem (NLP), that can be efficiently solved. Our proposed method is compared to previous methods in the literature that do not enforce aggregate cycling. Enforcing individual cycling constraint without taking that into account in reference generation leads to poor reference tracking.

I. INTRODUCTION

Successful and reliable operation of the power grid requires supply and demand to be balanced. In an envisioned future where renewable energy sources provide a higher percentage of energy, the imbalance becomes worse due to the inherent volatility of renewable sources. Additional ancillary services will be needed to maintain demand supply balance. Using fossil fuel based generators to provide ancillary services counteracts the benefits of renewables, while using batteries to do so is expensive [1]. One promising alternative to fossil fuel based sources and batteries is flexible loads. Flexible loads are consumers of energy that have the ability to either increase or decrease consumption over a baseline consumption level without affecting their Quality of Service (QoS).

In practice, flexible loads have been providing ancillary services for a long time under the domain specific term “demand response” [2], [3]. Demand response refers to the temporary reduction in demand, for a given load, to help

assist the grid. However, flexible loads have the potential to do much more. For instance, the demand of a flexible load could be automatically adjusted in real time so to track an exogenous power deviation (from baseline) reference signal. Baseline refers to the power consumption that would have occurred without interference from the grid operator. The reference signal will be provided by a balancing authority (BA). The tracking of this reference mimics that of a traditional storage service discharging and charging energy, so we refer to deviation of flexible loads’ power consumption from baseline as Virtual Energy Storage (VES) [4].

Some examples of flexible loads that are well studied in the academic literature are Thermostatically Controlled Loads (TCLs), commercial buildings, and pool pumps [5], [6], [7], [8], [9], [10]. In all of the above examples, ancillary services are possible because of “loose” requirements specified by baseline operation. For example, a residential Air Conditioner (an example of a TCL) is only required to maintain a temperature within a range and not at a specific value and a pool pump only desires an interval amount of ‘on time’ per day and not at a specific time. In this work, we focus on Air Conditioner (AC) TCLs.

In order for an aggregation of TCLs to track a reference signal, a model for the aggregate power consumption is required. A significant amount of research has been undertaken in developing models of aggregate power consumption, specifically models with controllable inputs [11]. Since if the input is “controllable” the complex coordination problem associated with scheduling aggregations of TCLs can be abstracted to that of control design. A large portion of these models rely on “probabilistically switching” individual TCLs [10], [8], [7], [12], [13], that is each TCL switches on or off with a certain probability. Regardless of the approach, these redesigns assume that the local thermostat controller will be replaced with a new controller that resides either locally at the TCL or at the BA.

Once a model and coordination method for a collection of TCLs have been developed, one is still pressed with the task of determining exactly what kind of reference signals the collection can track. In other words, what is the capacity of the collection? The capacity is determined by the QoS constraints of the individual TCLs, which include the user’s thermal comfort (temperature), her energy bill, and compressor cycling. A considerable amount of research has been done in characterizing capacity, but ignores cycling. In Hao et al [14], power and thermal energy constraints for the aggregate power deviation signal are developed by

[‡] corresponding author.

AC and PB are with the Dept. of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL 32601, USA. AB is with Inria Paris and also with DI ENS, École Normale Supérieure, CNRS, PSL Research University, Paris, France. The research reported here has been partially supported by the NSF through award 1646229 (CPS-ECCS), and the French National Research Agency grant ANR-16-CE05-0008. Additionally, we thank Naren Raman and Bo Chen for help with Casadi.

making analogs to an actual battery. Other authors have taken geometric approaches to model aggregate capacity [15].

While coordination methods proposed in previous works have considered local temperature QoS, the methods used to ensure local cycling QoS are ad-hoc (such as the so-called priority stack controller [14]). Applying them may satisfy local QoS at the cost of poor reference tracking (the BA's QoS). On the other hand, tracking a reference that exceeds the capacity of the collection - that is determined by all local QoS metrics - will require relaxing the local QoS.

Other works that investigate aggregate cycling constraints are [16], [17]. In [16] an aggregate cycling constraint is posed but reference generation from these constraints is not provided. That is, reference signals were constructed manually to test validity of the aggregate cycling constraints, but were not generated automatically. In [17], a centralized approach to handling aggregate cycling constraints is taken.

In this work we consider an *aggregate cycling capacity constraint*. The novel aspect of our work is twofold. First, an aggregate cycling constraint is derived. Second, we utilize this constraint to design a reference signal (in advance) that is close to the BA supplied reference signal, r_k^{BA} , that the TCLs can collectively track without violating their cycling QoS. In effect, the proposed method enables the BA to determine the reference closest to what the grid needs that the collection has sufficient capacity to deliver. Once the reference signal is generated, we borrow from our past work a distributed control framework that enables the reference signal to be tracked. We term our overall method as Cycling Aware Reference Generation and Tracking (CARG-T).

Our approach is different to those in [16], [17] in that the reference is generated automatically *and* the coordination algorithm is completely distributed. Additionally, the information requirement to generate the reference signal is not greater than previous reference generation methods that do not include aggregate cycling as a constraint [18].

We assume that the local TCL thermostat controller is replaced with a randomized controller from our past work [8], [9], [19]. This local controller induces (for each TCL) a new dynamic system, with a grid supplied signal as input and probability of being *on* as output. For a collection of TCLs, all listening to the same grid supplied signal and operating under the same local controller, a mean field limit exists. This limit allows for the individual dynamic system to be used as the aggregate model and the probability of any TCL being on to be interpreted as the fraction of total TCLs on. The key insight is that the aggregate model allows for a natural specification of our proposed aggregate cycling constraint directly in terms of the state variable of the model. The aggregate model is also shown to satisfy the aggregate power and thermal energy capacity constraints of [14]. Coupling of the aggregate model and aggregate cycling constraint allows for a description of the aggregate constraint set. This constraint set is then used to formulate the reference generation problem as a constrained optimization problem.

Tracking the generated reference with our distributed control architecture requires intelligence at both the grid level

(Balancing Authority, BA) and local level (individual TCL). The BA is responsible for solving the reference generation optimization problem, as well as determining a suitable control strategy to track the designed reference. The individual TCL is responsible for enforcing its own local QoS, which in this work are (i) temperature and (ii) device cycling rate. The randomized control policy mentioned above enforces the temperature QoS by design. A local opt-out control is coupled with the randomized controller to enforce local cycling QoS. Since the reference is generated by taking the capacity limitation of the collection due to cycling constraints of the individuals into account, the grid level and local level controllers are not in conflict. This enables reference tracking without violating local QoS.

To illustrate the overall effectiveness of our method, we compare our reference design and tracking method to a completely centralized priority stack controller [14]. The priority stack controller is a full state feedback controller located at the BA, that subjects the TCLs under its control to maintain all relevant local QoS requirements (including cycling). The control objective of the priority stack controller is to track a reference signal designed by the method presented in [18], which does not consider aggregate cycling capacity. It is then shown that reference tracking performance is poor when the priority stack controller attempts to track a reference that was not designed with aggregate cycling capacity in mind.

The rest of the paper is organized as follows: in Section II the grids needs are discussed, Section III describes the local randomized controller, Section IV and V present our aggregate model and proposed method, Section VII presents the comparison method, and Section VIII presents comparison results.

II. NEEDS OF THE GRID

It is envisioned that collections of TCLs will be a resource for the BA to help eliminate demand supply mismatch. In order to enable this, the BA must supply a reference signal, which we term r_k^{BA} . The BA-supplied reference is the desired deviation of the TCLs' power consumption from their baseline consumption. For instance, the Area Control Error (ACE) signal or its filtered versions (such as PJM's regD and regA [20]) can serve as the BA supplied reference.

As discussed in Section I, a collection of TCLs have limited capacity dictated by local QoS constraints. Therefore they may not be able to deliver the entirety of r_k^{BA} . We argue that the BA should therefore determine the portion of r_k^{BA} that the collection of TCLs could feasibly track. We term the feasible portion as r_k , and this is the signal the BA expects the collection to track. By design this reference signal will always be a zero mean deviation signal so that the TCLs are only asked to act as virtual batteries, but not asked to act as generators. At time k , the collection of TCLs need to consume r_k more/less than their baseline power consumption. Baseline power consumption refers to the power consumption that would have occurred without the interference from the BA, which we denote by \bar{P} (for the collection).

III. THE INDIVIDUAL TCL

A. QoS constraints

An on/off TCL is any device that turns on or off to maintain a temperature within a preset temperature deadband. Here, we denote the state space of a TCL as X , and elements of X as the couple $x = \{x^m, x^\theta\} = \{m \in \{0, 1\}, \theta \in \mathbb{R}\}$, that consists of the on/off status (mode) and temperature of the TCL. We denote the electrical power consumption as P , which is the power consumed by the TCL when it is on. Furthermore, time is discrete and denoted by the index k . The TCL index is j . For convenience, we also define the two sets $X^{\text{on}} \triangleq \{x : x^m = 1\}$ and $X^{\text{off}} \triangleq \{x : x^m = 0\}$. With the state variables declared, operating constraints (QoS) for the individual TCL are described as:

$$|\theta_k - \theta_{set}| \leq \delta, \quad \forall k \quad (1)$$

$$\sum_{i=0}^{\tau-1} |m_{k-i} - m_{k-1-i}| \leq 1, \quad \forall k \quad (2)$$

$$\left| \sum_{i=0}^{N_t-1} m_i P - \bar{E} \right| \leq \tilde{E}, \quad (3)$$

where \bar{E} is baseline energy consumption, and \tilde{E} is the permitted energy deviation during a time horizon N_t . The variables θ_{set} (set point) and δ (half width) specify the temperature deadband: $[\theta_{min}, \theta_{max}] = [\theta_{set} - \delta, \theta_{set} + \delta]$.

The first constraint is that the temperature remain within the temperature deadband. The second constraint is that the device can only switch once within a specified period τ , and the third constraint is that the individual device not consume significantly more (or less) energy (\tilde{E}) than baseline (\bar{E}). For instance, if $N_t = 1$ month then (3) represents a constraint to keep the energy bill of the consumer close to nominal.

B. Local Deterministic Control (Thermostat)

The local thermostat controller deterministically turns an AC TCL on when it is too hot and off when it is too cold, so to keep the temperature within the temperature deadband (1). We model the temporal evolution of θ_k as a first order discrete time equation

$$\theta_{k+1} = \bar{a}\theta_k + (1 - \bar{a})(w_{0,k} - Rm_k Q_{ac}), \quad (4)$$

$$\bar{a} = \exp\left(\frac{-T_s}{RC}\right), \text{ and } w_{0,k} = \theta_k^a, \quad (5)$$

where R and C represent the thermal resistance to ambient temperature (θ_k^a) and thermal capacitance of θ_k , respectively, and T_s is the sampling time. The thermal power consumption Q_{ac} is related to the electrical power consumption by $Q_{ac} = \eta P$, where η is the Coefficient of Performance (COP).

C. Local Randomized Control with Cycling Constraint

In this work each TCL determines its on/off status through a randomized rule [12], [8]. This randomized rule is a mapping from the state space of the TCL to an on/off (1/0) decision. In contrast to local deterministic control, where the

TCL can only turn on/off at $\theta_{max}/\theta_{min}$, in local randomized control the TCL can turn on/off at any temperature in $[\theta_{min}, \theta_{max}]$. However, the local randomized controller is specified so that a TCL will turn off/on at $\{\theta_{min}, \theta_{max}\}$, respectively. We now briefly review the randomized controller (see [19] for details).

The randomized controller is designed in two steps: (i) the nominal policy and (ii) the grid controlled policy. Here the nominal policy is specified through two random variables $\theta^{on}, \theta^{off}$ and their cumulative distribution functions (CDF), $F_{\Theta^{on}}(\theta), F_{\Theta^{off}}(\theta)$ that are supported on the temperature deadband of the TCL. The r.v.s $\theta^{on}(\theta^{off})$ are the temperatures at which the TCL switches on (off) after having started in the off (on) state, respectively. In this work the CDFs are chosen as:

$$F_{\Theta^{on}}(\theta) = \exp(-(\theta_{max} - \theta)^\rho / (2\sigma^\rho))$$

$$F_{\Theta^{off}}(\theta) = 1 - F_{\Theta^{on}}(\theta_{min} + \theta_{max} - \theta)$$

where ρ and σ are design parameters, and the values of the CDF are appropriately 1 or 0 outside the temperature deadband to ensure the essential CDF properties. In this paper we use $\rho = 0.75$ and $\sigma = 0.02$.

From the CDFs, the switching probability functions are defined as,

$$p_0^{on}(k) = \frac{[F_{\Theta^{on}}(\theta_k) - F_{\Theta^{on}}(\theta_{k-1})]_+}{1 - F_{\Theta^{on}}(\theta_{k-1})}$$

$$p_0^{off}(k) = \frac{[F_{\Theta^{off}}(\theta_{k-1}) - F_{\Theta^{off}}(\theta_k)]_+}{F_{\Theta^{off}}(\theta_k)}$$

where $[x]_+ = \max(x, 0)$, to ensure positive switching probabilities. These switching functions can be written compactly in the form of the Markov operator,

$$R_0(x, 1) = \begin{cases} p_0^{on}(k), & x^m = 0. \\ 1 - p_0^{off}(k), & x^m = 1. \end{cases} \quad (6)$$

$$R_0(x, 0) = \begin{cases} p_0^{off}(k), & x^m = 1. \\ 1 - p_0^{on}(k), & x^m = 0. \end{cases} \quad (7)$$

The second step in designing the randomized controller is the grid controlled policy. This enables the balancing authority to alter the switching probabilities through a control signal, that is broadcasted to every TCL, ζ . How this exactly happens is another design choice, and in this work the myopic policy of [21] is chosen. That is,

$$R_\zeta(x, y^m) = R_0(x, y^m) \exp(\zeta \mathcal{U}(y^m) - \Lambda_\zeta(x)) \quad (8)$$

where $\mathcal{U}(\cdot)$ is the *utility function*,

$$\mathcal{U}(x) = \mathcal{U}(x^m) = x^m, \quad (9)$$

and $\Lambda_\zeta(x)$ is a normalization constant to ensure probabilities sum to 1.

This translates to switching probability functions for when $\zeta \neq 0$,

$$p_\zeta^{on}(k) = \frac{p_0^{on}(k) \exp(\zeta)}{\exp(\Lambda_\zeta(x))} \quad (10)$$

$$p_\zeta^{off}(k) = \frac{p_0^{off}(k)}{\exp(\Lambda_\zeta(x))} \quad (11)$$

The temperature dynamics of the TCL are also written as the Markov operator, Q_{w_0} , where an element of Q_{w_0} , $Q_{w_0}(x, y^\theta)$ represents the transition probability to temperature state y^θ from state x due to disturbance w_0 . Construction of $Q_{w_0}(x, y^\theta)$ can either be from the specification of an analytical model, such as (4), or estimated through collected operational data of a TCL. In this work Q_{w_0} and R_ζ act independently, which allows us to write the transition from one state to another as the pointwise product,

$$P_{\zeta, w_0}(x, y) = R_\zeta(x, y^m)Q_{w_0}(x, y^\theta) \quad \forall x, y \in X. \quad (12)$$

C.1: Opt-out Control: An additional layer of local control is added to enforce the individual cycling constraint (2). This local control strategy is borrowed from [12], [22], which is termed ‘‘opt-out’’ control in [22]. The strategy is: once a device changes mode it must remain in that mode for τ sample times, where τ is as shown in (2).

IV. THE AGGREGATE MODEL

In order to control the aggregate power consumption of a collection of TCLs a model for the aggregate power consumption is required, which we term the aggregate model. We construct such an aggregate model of a collection of TCLs in which each operates according to the randomized controller described in Section III-C. The input to the aggregate model is ζ , the same input to individual TCLs. The outputs include the aggregate power consumption and fraction of mode switches. Opt-out control is not included in the aggregate dynamics as we seek necessary conditions to impose at the aggregate level so as to satisfy the individual QoS, and opt-out control represents a sufficient condition for individual QoS.

A. Aggregate TCL Model [21]

Let μ_k be the probability density function (pdf) (defined over the state space X) of the state of a TCL when it operates according to the randomized control law described in Section III-C. The pdf, μ_k , and the probability, γ_k , of the state being in the ‘‘on’’ mode, at time k , evolve according to

$$\mu_{k+1} = \mu_k P_{\zeta_k, w_0} \quad (13)$$

$$\gamma_k := P(x_k^m = 1) = \int_{X_{\text{on}}} \mu_k(x) dx. \quad (14)$$

As a system with inputs ζ, w and output γ , (13) is an infinite dimensional dynamic system that is linear in the state but nonlinear in the inputs. To make the connection between the individual and the aggregate, consider N homogeneous loads operating under the same randomized policy and subject to the same inputs (ζ_k, w_0) . Define the fraction of loads that are on at time k :

$$y_k^{(N)} := \frac{1}{N} \sum_{j=1}^N \mathcal{U}(x_k^{(j)}). \quad (15)$$

A mean-filed limit holds, proven in [19]. That is, if N homogeneous loads are subjected to the same inputs ζ and

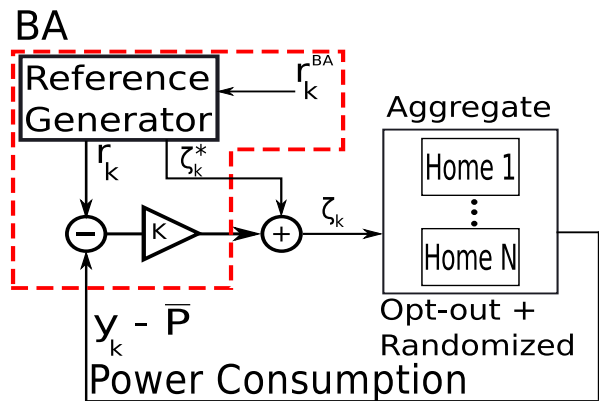


Fig. 1. Distributed load control architecture of the proposed method (CARG-T).

w , as $N \rightarrow \infty$ the fraction of loads that are on at time k approaches the probability of a single load being on at k :

$$\lim_{N \rightarrow \infty} y_k^{(N)} = \gamma_k \quad (16)$$

In the sequel, for simplicity we drop the superscript N from all quantities that contain an average over N .

This allows us to treat (13) as the state equation for the population. The relevant output equations, in this work, can now be computed through the density $\mu_k(x)$,

$$y_k = P_{\text{agg}} \int_{X_{\text{on}}} \mu_k(x) dx, \quad P_{\text{agg}} := \sum_{j=1}^N P_j, \quad \text{and} \quad (17)$$

$$\bar{s}_k = \int_{X_{\text{on}}} p_{\zeta_k}^{\text{off}}(x) \mu_k(x) dx + \int_{X_{\text{off}}} p_{\zeta_k}^{\text{on}}(x) \mu_k(x) dx. \quad (18)$$

Here \bar{s}_k is the fraction of switches (normalized by the number of TCLs) from timestep k to $k+1$ and y_k is the power consumption of the collection. Note that except (18) both (13) and (17) are borrowed from [19]. The eq. (18) provides a way to compute the number of switches in terms of the state variable $\mu_k(x)$.

V. PROPOSED METHOD: CARG-T

Here we present our proposed method, which we term Cycling Aware Reference Generation and Tracking (CARG-T). The novel aspect of our method is twofold: (i) the construction of an aggregate cycling constraint and (ii) utilization of this constraint to design arbitrary time varying reference signals (in advance) that aggregations of TCLs can track without violating their individual cycling QoS. The aggregate cycling constraint acts as a necessary condition for the individual TCLs cycling constraint (2). That is, the reference is designed so that it can be successfully tracked, even when the TCLs enforce the local cycling QoS. Solving the reference generation problem also produces an open-loop control sequence, $\{\zeta_k^*\}$. To implement the reference we couple this open loop sequence with a proportional feedback controller. An illustration of the overall architecture is shown in Figure 1.

A. Aggregate Cycling Constraint

We determine a cycling constraint for the aggregate that is consistent with the cycling constraint for the individual TCL. This constraint is obtained by assuming homogeneous individual cycling constraints and summing the individual cycling constraint (2) over all of the devices:

$$\sum_{j=1}^N \sum_{i=0}^{\tau-1} s_{k-i}^j \leq N, \quad s_k^j := |m_k^j - m_{k-1}^j| \quad (19)$$

$$\sum_{i=0}^{\tau-1} \bar{s}_{k-i} \leq 1, \quad \text{where } \bar{s}_{k-i} := \frac{1}{N} \sum_{j=1}^N s_{k-i}^j. \quad (20)$$

In (19), \bar{s}_k is the fraction of TCLs that switch from time k to $k+1$. The aggregate cycling constraint (20) is a necessary condition for the individual cycling constraints at the TCLs to hold, as the aggregate fraction of switches will only exceed 1 if any subset of devices switch more than once in the past τ time.

The fraction of switches can be estimated by summing, over all of the TCLs, the probability of switching for each TCL and then dividing by the number of TCLs. Utilizing the density $\mu_k(x)$ and probability switching function $p_k(x) := [p_{\zeta_k}^{on}(x), p_{\zeta_k}^{off}(x)]^T$ for the TCLs, this is expressed as,

$$\bar{s}_k = \int_X \mu_k(x) p_k(x) dx. \quad (21)$$

This is equivalent to the output equation for the fraction of switches shown in (18).

B. Reference Generation

We couple the mean field model (13) with the aggregate cycling constraint (20) to form an aggregate constraint set. We utilize the aggregate constraint set to solve an optimization problem to generate a power deviation reference trajectory that satisfies this constraint set:

$$\min_{\{\zeta_k\}_{k=0}^{N_t-1}, \{\mu_k\}_{k=1}^{N_t}} \sum_{k=1}^{N_t} (r_k - r_k^{BA})^2 \quad (22)$$

s.t. $\forall k \in \{0, \dots, N_t\}$

$$r_k = P_{agg} \int_{X_{on}} \mu_k(x) dx - \bar{P} \quad (23)$$

$$\sum_{i=0}^{\tau-1} \bar{s}_{k-i} \leq 1, \quad \bar{s}_k = \int_X \mu_k(x) p_k(x) dx \quad (24)$$

$$\mu_{k+1} = \mu_k P_{\zeta_k, w_0}, \quad \mu_0 = \mu \quad (25)$$

$$\frac{1}{N_t} \sum_{k=1}^{N_t} r_k = 0. \quad (26)$$

The problem (22) is implemented on the computer through discretization of the state space of (13), which turns (22) into a finite dimension non-linear optimization problem.

The constraint (26) ensures (3) on average, with $\tilde{E} = 0$. This constraint can be interpreted as a necessary condition for the individual energy deviation QoS constraint (3). For if this constraint were not imposed the collection of TCLs

could consume more or less (in mean) than they should have, which would correspond to higher electricity bills for some home owners.

Here we have used $\tau_p = 2\tau$, as a factor of safety in reference design. The baseline \bar{P} is precomputed from a simulation of TCLs providing no ancillary services. Additionally, r_k^{BA} is a forecasted signal that the BA would like the TCLs to track (see Section II).

The goal in solving this problem is to determine r_k^* , the closest feasible power reference trajectory to r_k^{BA} , which is what the BA needs, that respects all of the individual TCLs QoS constraints.

C. Reference Tracking

The solution to (22) also yields an open loop control sequence $\{\zeta_k^*\}_{k=0}^{N_t-1}$. This control sequence is then coupled with a proportional feedback controller,

$$\zeta_k = \zeta_k^* + \frac{K_p}{P_{agg}} (r_k^* - (y_k - \bar{P})) \quad (27)$$

where y_k is the aggregate power consumption for the collection of TCLs during implementation. The above control sequence can be used by the BA as a controller, and broadcast ζ_k to all TCLs. Each TCL implements the randomized controller with opt-out control, as described in Section III-C. Opt-out will ensure that (2) is satisfied for each TCL and recall, by design the randomized controller ensures that each TCL will not violate (1).

VI. AGGREGATE POWER AND THERMAL ENERGY CONSTRAINTS

In He Hao et al [14], aggregate power and *thermal* energy constraints were proposed. These constraints arise by making a comparison between (4) and the ODE model of a “leaky” battery,

$$z_{k+1} = \bar{a}z_k - br_k, \quad b = (1 - \bar{a})RC \quad (28)$$

$$|z_k| \leq \bar{C} \quad \eta_- \leq r_k \leq \eta_+. \quad (29)$$

Here z_k represents the thermal state of charge for a collection of TCLs, and can be interpreted as the average (over space) temperature deviation of the collection of TCLs from the setpoint of each TCL. r_k is the reference signal for the battery (desired charge or discharge). The parameter \bar{a} is as given in (5). For a homogeneous collection of TCLs, the parameters $\{\bar{C}, \eta_-, \eta_+\}$ are [14]:

$$\bar{C} = \frac{NC\Delta}{2\eta}, \quad (30)$$

$$\eta_- = -\bar{P}, \quad \eta_+ = P_{agg} - \bar{P}, \quad (31)$$

where $\Delta = 2\delta$, i.e. the full width of the temperature dead-band. The inequality constraint values $\{\bar{C}, \eta_-, \eta_+\}$ define the aggregate power and thermal energy capacity of the collection.

In this work, we show that the mean field model (13) (and thus the CARG-T reference generator) satisfies these constraints for a homogeneous population.




Fig. 2. Centralized Direct load control architecture of the comparison method (CHRG-T).

The energy constraint is shown by considering the quantity, $\phi_k(x) \triangleq (NC\mu_k(x)/\eta)(x_k^\theta - \theta_{set})$. This quantity, $\phi_k(x)$, is a representation of z_k in (28) when it is integrated over the state space (i.e., the average temperature deviation from θ_{set}). From this, it can be shown that the thermal energy constraint (29) is satisfied by integrating $\phi_k(x)$ over X ,

$$\left| \int_X \phi_k(x) dx \right| = \frac{NC}{\eta} \left| \int_X x_k^\theta \mu_k(x) dx - \theta_{set} \right| \quad (32)$$

$$\leq \frac{NC}{\eta} |(\delta + \theta_{set}) - \theta_{set}| = \frac{NC\delta}{\eta} = \frac{NC\Delta}{2\eta} = \bar{C}. \quad (33)$$

With the integral of $\phi_k(x)$ playing the role of z_k , (29) (left) is satisfied for any realization of the density $\mu_k(x)$ so no thermal energy constraint is forcibly required for (13).

To show the power constraint consider,

$$0 \leq P_{agg} \left| \int_X \mu_k(x) dx \right| \leq P_{agg} \quad (34)$$

$$-\bar{P} \leq P_{agg} \left| \int_X \mu_k(x) dx \right| - \bar{P} \leq P_{agg} - \bar{P} \quad (35)$$

since $\mu_k(x)$ is a pdf. So (29) is satisfied (see (31)).

VII. COMPARISON METHOD: CHRG-T [18], [14]

Here we present an alternative method for generating and tracking a reference trajectory for a collection of TCLs that does not take care of aggregate cycling. The reference generation portion is taken from [18] and the tracking portion is taken from [14]. The reference generation problem considers aggregate power and thermal energy constraints (29). The reference tracking portion is a centralized priority stack controller, meaning all intelligent decisions are made at the BA. We term this alternative reference generation and tracking approach as Cammardella + Hao et al. - Reference Generation and Tracking (CHRG-T), and offer it as a comparison method to our proposed algorithm, CARG-T. The architecture for this proposed method is illustrated in Figure 2.

A. Reference Generation

First we detail the reference generation portion from [18], and we re-present it here:

$$\min_{\{r_k\}_{k=0}^{N_t-1}, \{z_k\}_{k=1}^{N_t}} \sum_{k=0}^{N_t-1} (r_k - r_k^{BA})^2 \quad (36)$$

$$\text{s.t. } \forall k \in \{0, \dots, N_t - 1\}$$

$$z_{k+1} = \bar{a}z_k - br_k, \quad z_0 = z \quad (37)$$

$$|z_k| \leq \bar{C}, \quad \eta_- \leq r_k \leq \eta_+ \quad (38)$$

$$\frac{1}{N_t} \sum_{k=0}^{N_t-1} r_k = 0. \quad (39)$$

Here, r_k represents the reference power deviation of the collection, z_k represents the thermal state of charge of the collection, and r_k^{BA} is the same signal as in (22). This optimization problem considers the aggregate power and thermal energy constraints presented in [14], as well as the aggregate energy deviation constraint (39).

Comment 1: Both optimization problems (22) and (36) require the same amount of information about the population of TCLs in order to be set up and solved.

TABLE I
PARAMETERS

Classification	Symbol	Value	Units
BA parameters			
Number of TCL's	N	60000	N/A
Thermal Energy bound	\bar{C}	53	MWh
Upper Power deviation bound	η_+	35	MW
Lower Power deviation bound	η_-	-65	MW
Baseline Consumption	\bar{P}	65	MW
Planning Opt-out Time	τ_p	40	minutes
Feedback Gain	K_p	10	N/A
TCL parameters			
Resistance for TCL model	R	$U[2, 2.25]^*$	K/kW
Capacitance for TCL model	C	$U[2, 2.25]$	kWh/K
Rated elec. power	P	$U[1.6, 1.8]$	kW
Coeff. of Perf.	η	2.5	NA
Temperature set point	θ_{set}	21	$^\circ C$
Temperature deadband	δ	1	$^\circ C$
Ambient Temperature	θ^a	30	$^\circ C$
Device Opt-out Time	τ	20	minutes

* $U[a, b]$ is uniformly distributed between a and b.

B. Reference Tracking

Since the authors in [18] do not provide a control strategy to track the *generated* reference signals, we adopt the control algorithm developed in [14]. The reason for this is that the capacity model utilized in CHRG-T originates from [14]. The control algorithm in [14] is a *direct centralized load control* (DLC) framework for controlling aggregate power consumption of a collection of TCLs. DLC refers to the fact that the on/off status of the TCLs are controlled directly. At the BA is a priority stack controller that determines the on/off status of each TCL. This controller *requires full state information* (temperature and on off status for every TCL in the population) in order to determine which TCL's mode should be switched. In this work, we augment the priority




Fig. 3. The original balancing reserves signal, r_k^{BA} , and the two designed reference signals.




Fig. 5. Reference tracking results (60000 TCLs) with CARG-T.




Fig. 4. Individual cycling QoS (60000 TCLs) with CARG-T.

stack controller with opt-out, so that device's local cycling QoS constraint is satisfied. Since all of the decision making is done at the BA, there is no need for a local controller at each TCL. Since the priority stack controller has full state information, it has an inherent unfair advantage over our distributed controller.

VIII. RESULTS

We present numerical results of our proposed method, CARG-T, and the comparison method, CHRGT-T. For both methods, data for r_k^{BA} is obtained from Bonneville Power Administration (BPA), a BA in the Pacific northwest. Data for the parameters of the individual TCL's are described in Table I, which are obtained from [23]. The ambient temperature, $30^\circ C$, is uniform over all the TCL's and constant over time, which translates to a constant baseline value. In our proposed method CARG-T, the reference generation problem (22) is solved using Casadi [24] and the NLP solver IPOPT [25]. In the comparison method CHRGT-T, the reference generation problem (36) is solved with CVX [26].

TABLE II
REFERENCE TRACKING ERRORS

Method	Tracking Error
Proposed Method (CARG-T)	2.01 %
Comparison Method (CHRGT-T)	15.67 %




Fig. 6. Individual cycling QoS results (60000 TCLs) with CHRGT-T.

A. Reference generated

In Figure 3, the original balancing reserves signal, r_k^{BA} , and the two designed reference signals are shown. Both reference signals share a similar trend, although it is clear that CHRGT-T is more aggressive compared to CARG-T. That is, the reference signal, designed with CHRGT-T, attempts to track more of r_k^{BA} , with higher ramp rates than that by CARG-T. Since CARG-T enforces additional aggregate cycling constraints, this result is expected.

B. Reference Tracking and QoS

For the proposed method CARG-T the individual cycling QoS results are shown in Figure 4 and the reference tracking results are shown in Figure 5. For the comparison method CHRGT-T the individual cycling QoS results are shown in Figure 6 and reference tracking results are shown in Figure 7. A comparison of the numerical reference tracking error is also shown in Table II. Recall, CHRGT-T does not enforce the aggregate QoS cycling constraint in generating a reference and it can be observed that reference tracking performance suffers (Figure 7 and Table II). This poor performance is because of the opt-out control, which is a *requirement* for each TCL to maintain successful operation. Removal of opt-out control would improve reference tracking, but at the cost of devices violating their cycling QoS.

Also, the time between device switches (shown in Figure 6) is highly synchronized in the comparison CHRGT-T method. This is in contrast to the proposed CARG-T where the time between successive switches is evenly spread out. The poor reference tracking results for CHRGT-T is not the fault of the controller, rather, due to the reference not being




Fig. 7. Reference tracking results (60000 TCLs) with CHR-G-T.

generated with cycling constraints in mind. It should be noted that both methods enforce the local temperature QoS by design.

IX. CONCLUSION

The proposed method enables the BA to determine the reference closest to what the grid needs that a collection of TCLs has sufficient capacity to deliver. Once the reference signal is generated, we borrow from our past work a distributed control framework that enables the reference signal to be tracked. The proposed method is compared against a *centralized* controller and reference generator that does not enforce aggregate cycling. The centralized controller, by design, enforces the local cycling constraint. It is seen that this centralized controller is unable to track the reference and leads to synchronized cycling of TCLs. In contrast our proposed method is able to track the reference it generates without synchronizing the TCLs, and keep the cycling rate of the individual within allowable limits. A key future work is to develop aggregate level sufficient conditions for individual level QoS constraints.

REFERENCES

- [1] N. J. Cammardella, R. W. Moye, Y. Chen, and S. P. Meyn, "An energy storage cost comparison: Li-ion batteries vs distributed load control," in *2018 Clemson University Power Systems Conference (PSC)*, Sep. 2018, pp. 1–6.
- [2] Y. V. Makarov, J. Ma, S. Lu, and T. B. Nguyen, "Assessing the value of regulation resources based on their time response characteristics," *Pacific Northwest National Laboratory*, June 2008.
- [3] P. Barooah, A. Bušić, and S. Meyn, "Spectral decomposition of demand side flexibility for reliable ancillary service in a smart grid," in *48th Hawaii International Conference on Systems Science (HICSS)*, January 2015, invited paper.
- [4] P. Barooah, *Smart Grid Control: An Overview and Research Opportunities*. Springer Verlag, 2019, ch. Virtual energy storage from flexible loads: distributed control with QoS constraints, pp. 99–115.
- [5] S. Bashash and H. K. Fathy, "Modeling and control of aggregate air conditioning loads for robust renewable power management," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1318–1327, 2013.
- [6] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.
- [7] W. Zhang, J. Lian, C. Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, Nov 2013.

- [8] Y. Chen, M. U. Hashmi, J. Mathias, A. Bušić, and S. Meyn, "Distributed control design for balancing the grid using flexible loads," in *Energy Markets and Responsive Grids*. Springer, 2018, pp. 383–411.
- [9] A. R. Coffman, A. Busic, and P. Barooah, "A study of virtual energy storage from thermostatically controlled loads under time-varying weather conditions," in *5th International High Performance Buildings Conference at Purdue*, 2018.
- [10] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.
- [11] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.
- [12] A. Coffman, A. Busic, and P. Barooah, "Virtual energy storage from tcls using qos preserving local randomized control," *5th ACM international conference on systems for the built environment*, vol. 5, Jan 2018.
- [13] M. Liu and Y. Shi, "Model predictive control of aggregated heterogeneous second-order thermostatically controlled loads for ancillary services," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 1963–1971, May 2016.
- [14] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 189–198, Jan 2015.
- [15] L. Zhao, W. Zhang, H. Hao, and K. Kalsi, "A geometric approach to aggregate flexibility modeling of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 32, no. 6, pp. 4721–4731, Nov 2017.
- [16] C. Ziras, S. You, H. W. Bindner, and E. Vrettos, "A new method for handling lockout constraints on controlled tcl aggregations," in *2018 Power Systems Computation Conference (PSCC)*, June 2018, pp. 1–7.
- [17] B. M. Sanandaji, T. L. Vincent, and K. Poolla, "Ramping rate flexibility of residential hvac loads," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 865–874, April 2016.
- [18] N. Cammardella, J. Mathias, M. Kiener, A. Bušić, and S. Meyn, "Balancing California's grid without batteries," in *2018 IEEE Conference on Decision and Control (CDC)*, Dec 2018, pp. 7314–7321.
- [19] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, "Ancillary service to the grid from intelligent deferrable loads," *IEEE Transactions on Automatic Control*, vol. 60, pp. 2847–2862, March 2015.
- [20] PJM, "PJM manual 11: Energy and ancillary services market operations, rev. 62," December 2013.
- [21] A. Bušić and S. Meyn, "Distributed randomized control for demand dispatch," in *IEEE conference on decision and control*, 2016, pp. 6964–6971.
- [22] Y. Chen, A. Bušić, and S. Meyn, "Individual risk in mean field control with application to automated demand response," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 6425–6432.
- [23] J. L. Mathieu, M. Dyson, and D. S. Callaway, "Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations," in *In Proceedings of the ACEEE Summer Study on Buildings*, 2012.
- [24] J. Andersson, J. Åkesson, and M. Diehl, "Casadi: A symbolic package for automatic differentiation and optimal control," in *Recent Advances in Algorithmic Differentiation*, S. Forth, P. Hovland, E. Phipps, J. Utke, and A. Walther, Eds. Springer Berlin Heidelberg, 2012, pp. 297–307.
- [25] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, Mar 2006. [Online]. Available: <https://doi.org/10.1007/s10107-004-0559-y>
- [26] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," <http://cvxr.com/cvx>, Mar. 2014.