

Managing Voltage Excursions on the Distribution Network by Limiting the Aggregate Variability of Thermostatic Loads

Stephanie C. Ross, Petter Nilsson, Necmiye Ozay, and Johanna L. Mathieu

Abstract—This paper proposes a strategy to control a group of thermostatically controlled loads (TCLs) such that the variability in their aggregate load is reduced. This strategy could be deployed in areas of a distribution network that experience voltage excursions due to net load fluctuations, such as areas with high penetrations of photovoltaic (PV) generation and/or electric vehicles (EVs). We limit variation in the power consumption of a group of TCLs using a control strategy previously developed for large aggregations of switched systems. Using this strategy, we constrain the number of TCLs that are on (i.e., actively consuming power) between upper and lower bounds. In simulations, the control strategy successfully decreases the range over which TCL power consumption varies. Percent reductions in range are greatest for medium group sizes: we find a median reduction of 82% for groups of 50 TCLs, 74% for groups of 1000 TCLs, and 59% for groups of 5 TCLs. Reducing the variability of a distribution network’s power injections helps to reduce voltage variability. In a simulation of a distribution line supplying 25 households, half with PV systems, the control strategy reduces the total range of voltage by 0.02 p.u. and prevents a violation of the 0.95 p.u. limit. Lastly, we propose a new control strategy for a more realistic TCL model that includes compressor lockout. The new strategy performs comparably to the original strategy and is demonstrated through simulation.

I. INTRODUCTION

As distribution-level photovoltaic (PV) generation continues to grow, so too does the possibility of large, fast changes in net load. Aggregate PV power generation can drop by 60% in less than 30 seconds when a cloud passes over a neighborhood [1]. There is also substantial variation in the aggregate power consumption of small groups of thermostatically controlled loads (TCLs) because of random periods of synchronization in their on/off power cycles. Residential TCLs, such as air conditioners (ACs) and water heaters, cycle on and off to regulate temperature. In an area of a network with high penetration of PV and TCLs, a simultaneous drop in PV generation and increase in TCL consumption could result in under-voltages, particularly if the systems are located at the end of a long distribution line.

Conventionally, distribution voltages are regulated with on-load tap-changing transformers, voltage regulators, and capacitor banks. However, these techniques have not been

designed to mitigate fast, repetitive excursions in voltage; for example, tap changers typically have built-in delays to avoid responding to transient conditions [2]. New forms of voltage management are needed in network areas with highly variable power injections and high voltage sensitivity to these variations (e.g., at nodes far from the substation with high TCL and PV penetration).

In this paper, we propose controlling groups of co-located TCLs to reduce the variation in their aggregate power consumption and thereby mitigate their contribution to local voltage excursions. The control strategy constrains the “on-count”, i.e., number of TCLs that are on simultaneously, between lower and upper bounds. We set the bounds as close together as possible so that the on-count, and resulting power consumption, are maximally constrained. There has been related work, in terms of reducing on/off synchronization of TCLs, for the prevention of power oscillations or rebound after a period of load reduction [3], [4]. However, these strategies control TCLs by adjusting temperature setpoints or ranges; in contrast, we control TCLs by switching them on/off, and do so within the user-set temperature range so that the control is non-disruptive to the end-user [5].

Prior work on non-disruptive control of TCLs has focused on aggregating hundreds to thousands of TCLs to provide transmission-level services (e.g., [6]–[8]), rather than providing distribution-level services. Distribution voltage constraints are considered within [9], [10], but the optimization-based algorithms would likely be too computationally intensive to respond to voltage excursions at the sub-minute timescale. Finally, the real-time algorithm of [11] incorporates distribution voltage measurements and constraints; however, the algorithm requires controllable resources to have continuously variable power outputs, rather than the discrete power outputs of residential TCLs.

The contributions of this paper are as follows:

- Through simulations, we demonstrate that large, fast voltage excursions could occur if variations in PV generation and aggregate TCL power consumption coincide.
- We apply the theory developed in [12], [13] on counting problems of switched systems to the real-world problem of distribution voltage management.
- We examine the effectiveness of the control strategy for different size groups of TCLs.
- Lastly, we extend the control strategy of [12] to the case in which TCLs have a “lockout period” after switching and cannot switch again until after the period has elapsed.

*This work was supported by the National Science Foundation under Grant Nos. CNS-1446298, CNS-1837680, DGE-1256260, ECCS-1553873, as well as the Rackham Predoctoral Fellowship.

S.C. Ross, N. Ozay, and J.L. Mathieu are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA {sjcrock, necmiye, jlmath}@umich.edu

P. Nilsson is with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125 USA pettni@caltech.edu

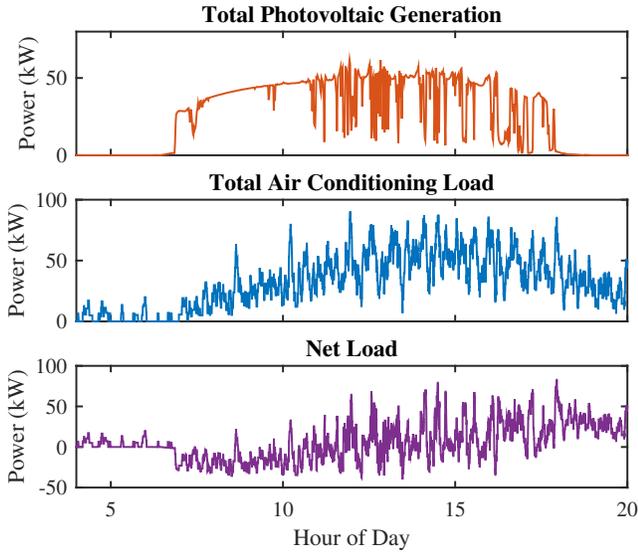


Fig. 1. Simulated power profiles of PV generation, AC consumption, and net load. The PV generation profile drops sharply due to partial clouding. Variation in the ACs' power consumption is due to random periods in which ACs are synchronized. The variation of the net load is larger than that of either resource alone.

II. VARIABILITY OF NET LOAD WITH HIGH PENETRATION OF PV AND TCLS

A. Variation in PV Power Output

If multiple PV systems are located in close proximity, changes in power output due to transient clouds will occur nearly simultaneously and could cause voltages to fluctuate in that area of the distribution network. We use the following model and simulation to demonstrate this effect.

1) *PV Model*: We use a simplified version of the model in [14] that estimates PV power output as a function of solar irradiance. Neglecting temperature effects, PV power output can be calculated as

$$P_{pv}(t) = \eta_{inv} D P_{dc0} S(t) / S_{ref}. \quad (1)$$

Parameters are defined in Table I with values sourced from [14], [15]. The variable S is the incident irradiance on the PV surface; we approximate the incident irradiance as the sum of direct normal irradiance and diffuse horizontal irradiance, which should overestimate PV production [14]. Lastly, we assume PV systems are controlled to inject power at a unity power factor.

2) *PV Simulation*: Using solar irradiance data from Oak Ridge National Lab [16], we simulate the total output power of 12 residential PV systems on a partially cloudy day (see Fig. 1 (top)). The irradiance data has 1 minute resolution and was measured in Oak Ridge, TN on Sept. 15, 2018 [16]. As shown in Fig. 1, partial clouds can cause PV generation to drop quickly; for example, at 12:42 pm the power output of the PV systems decreases by 82% in less than 1 minute.

B. Variation in TCL Power Consumption

A group of TCLs co-located at a distribution node can also cause large variations in power injections due to random

TABLE I
PHOTOVOLTAIC PARAMETERS

Parameter	Values	Unit
Inverter efficiency (η_{inv})	96%	–
Derating factor (D)	0.86	–
Rated DC power capacity (P_{dc0})	5.7	kW
Reference irradiance for rated power (S_{ref})	1000	W/m ²

periods of synchronization. We use the following model and simulation to demonstrate this effect.

1) *TCL Model*: We model a group of heterogeneous ACs using the thermal model developed in [17] and commonly used in the literature [18], [19]. The i th AC's indoor temperature θ_i evolves according to the hybrid dynamics,

$$\frac{d}{dt} \theta^i(t) = \begin{cases} -\frac{1}{\tau^i} (\theta^i(t) - \theta_a(t) + P_{\Theta}^i R^i) & \text{if } \sigma^i(t) = \text{on}, \\ -\frac{1}{\tau^i} (\theta^i(t) - \theta_a(t)) & \text{if } \sigma^i(t) = \text{off}, \end{cases} \quad (2)$$

where σ^i is the power mode, $\tau^i = C^i R^i$, and θ_a is the outdoor temperature. Parameters R , C , and P_{Θ} are defined in Table II; values for R and C are sampled from the random uniform distributions defined by the ranges in the table. The value for P_{Θ} is correlated with the value for R such that

$$\frac{P_{\Theta}^i - P_{\Theta, \min}}{P_{\Theta, \max} - P_{\Theta, \min}} = 1 - \frac{R^i - R_{\min}}{R_{\max} - R_{\min}}. \quad (3)$$

An AC's internal thermostat control determines the state of σ^i , when the AC is not subject to external control. The thermostat maintains the indoor temperature between the upper temperature limit $\bar{\theta}^i$ and lower limit $\underline{\theta}^i$ with the hysteretic control:

$$\sigma^i(t) = \begin{cases} \text{on} & \text{if } \theta^i(t) \geq \bar{\theta}^i, \\ \text{off} & \text{if } \theta^i(t) \leq \underline{\theta}^i, \\ \sigma^i(t^-) & \text{otherwise.} \end{cases} \quad (4)$$

where $\bar{\theta}^i = \theta_{set}^i + \delta_i/2$ and $\underline{\theta}^i = \theta_{set}^i - \delta_i/2$. The parameters θ_{set} and δ are defined in Table II and are drawn from the uniform distributions defined by the range of values in the table.

Finally, the real power consumption of the i th AC is

$$P_{AC}^i(t) = \mathbb{1}_{\{\text{on}\}}(\sigma^i(t)) P_{\Theta}^i / C_{cop}, \quad (5)$$

where C_{cop} is the coefficient of performance and is set to 2.5 for all ACs. We use the notation $\mathbb{1}_{\{a\}}(b)$ to represent the indicator function of the singleton set containing a : if $b = a$ the function has value 1 and if $b \neq a$ the function has value 0. We assume that when an AC is on, it is a constant power load with a power factor of 0.97. (Power factor values for residential loads are from Table A.2 of [20].)

2) *TCL Simulation*: We simulate the power consumption of a group of 25 heterogeneous ACs as a function of outdoor temperature (see Fig.1 (middle)). The outdoor temperature data has 1 minute resolution and is from the same day and location as the solar data [16]. Variations in power are of a similar magnitude and frequency as those of the PV

TABLE II
AIR CONDITIONER PARAMETERS

Parameter	Values	Unit
Setpoint temperature (θ_{set})	18-27	$^{\circ}\text{C}$
Width of temperature range (δ)	0.25-1	$^{\circ}\text{C}$
Thermal resistance (R)	1.2-2.5	$^{\circ}\text{C}/\text{kW}$
Thermal capacitance (C)	1.5-2.5	$\text{kWh}/^{\circ}\text{C}$
Thermal energy transfer rate (P_{Θ})	10-18	kW

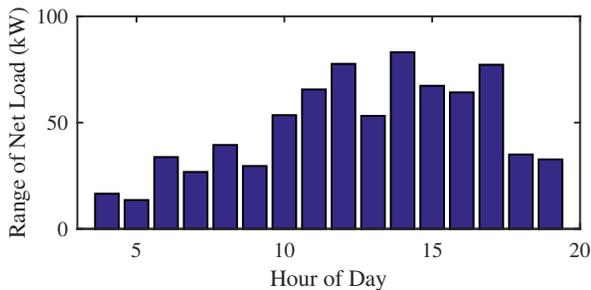


Fig. 2. Range of net load over each hour of the simulated day. The range is largest in the afternoon when PV generation and AC consumption both have large variations.

systems; for example, at 1:22 pm power consumption by ACs decreases by 89% in 6 minutes.

C. Net Load Variation

When groups of ACs and PVs are co-located at the same node, the variation of the net power injection can be greater than that of either resource independently. The bottom plot of Fig. 1 shows the net load profile, which is just the difference of the PV profile (top) and the AC profile (middle). The total range in power for the three profiles in Fig. 1 is 62 kW, 90 kW, and 122 kW for the PV, AC, and net load profiles, respectively. Figure 2 shows the range of the net load in each hour of the simulated day; the range is largest in the mid-afternoon when large variations in PV generation and AC consumption coincide. Note, in Fig. 2, range is calculated as the difference between the 5th and 95th percentiles.

D. Voltage Excursions

Large and fast variations in net load can cause voltage excursions at the end of a long distribution line. We use the following model and simulation to demonstrate this effect.

1) *Distribution System Model*: As shown in Fig. 3, we model a distribution line that connects two nodes: node 1, an infinite bus with voltage of 1.0 p.u., and node 2, a PQ bus that supplies 25 households. We assume twelve of these households also have a PV system. Each household has an AC and a constant power load that aggregates all other loads. The constant power load draws 1.5 kW of real power and has a power factor of 0.95. The PV systems and ACs are modeled as described in Sections II-A and II-B.

We use an intentionally simple distribution system model so that we can fully interpret the results. The line is single phase, and we omit distribution transformers between node 2 and the loads. Node 1 models a substation with voltage

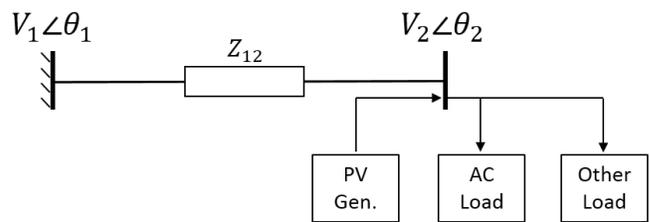


Fig. 3. Diagram of distribution line between infinite bus (node 1) and load bus (node 2). The line is single phase, 10 miles long, and supplies a group of 25 households.

regulation. Given these assumptions, we derive the following power flow equations:

$$P_{\text{net}} = V_1 V_2 (g_{12} \cos \theta_{21} + b_{12} \sin \theta_{21}) - g_{12} V_2^2 \quad (6)$$

$$Q_{\text{net}} = V_1 V_2 (g_{12} \sin \theta_{21} - b_{12} \cos \theta_{21}) + b_{12} V_2^2, \quad (7)$$

where P_{net} and Q_{net} are the net power injections at node 2, V_1 and V_2 are the nodes' voltage magnitudes, θ_{21} is the difference in voltage angle between the two nodes ($\theta_2 - \theta_1$), and parameters g_{12} and b_{12} are the line's conductance and susceptance, respectively. The line is 10 miles long and consists of a single phase and neutral, both of which have a 180-ampere conductor with parameters from a feeder model provided by [21]. We calculate the line's per mile impedance to be $1.86 + 1.41j$ using the "Modified Carson's Equations" (see Chapter 4 of [2]).

2) *Distribution System Simulation*: We simulate the power flow along the modeled distribution line using the same daily weather data as in previous sections. We solve for node 2's voltage in each time step by applying the Newton Raphson method to (6)-(7). The voltage profile of node 2 is shown in Fig. 4. In the middle of the day, large variations in the net power injection at node 2 cause voltage excursions below 0.95 p.u. (the national standard for service voltage is 0.95-1.05 p.u. [22]).

In a real system, the under-voltages in Fig. 4 could be prevented by increasing the setpoint of the voltage regulator at node 1 above 1 p.u. However, a higher setpoint could result in over-voltages on a sunny day with minimal load. Because of this trade-off, we expect that some systems' voltage regulation schemes will not be able to prevent all possible voltage excursions, and excursions similar to those in Fig. 4 will sometimes occur.

III. LIMITING THE VARIABILITY OF TCLs

We propose reducing the variability of TCL power consumption in order to reduce violations of voltage limits on distribution systems. This strategy is best suited for areas of the distribution system where voltage is highly sensitive to changes in power injections or where voltages are already close to their limits. Moreover, it is a preventive strategy, useful in situations when the load controller does not have real-time information about PV production or system voltages. For example, a utility might know that the weather forecast is for a partially cloudy day and limit the aggregate variability of TCLs in select areas as a preventive measure.

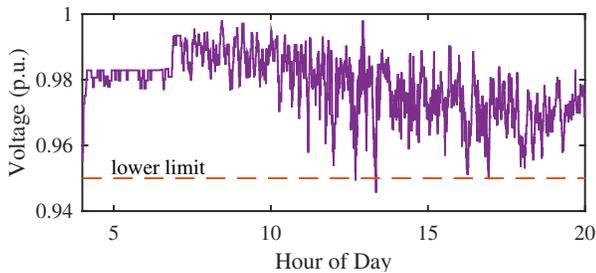


Fig. 4. Voltage profile at node 2 on the simulated distribution line. Voltage excursions reach below the 0.95 p.u. lower limit during large positive variations in net load.

A. Counting Problems for Aggregations of Switched Systems

We propose limiting the variability of a TCL group’s power consumption by constraining the number of TCLs that are on at any instant, i.e., the group’s “on-count”. We use the theory developed in [12], [13] on control strategies for aggregations of switched systems with counting constraints. In the remainder of Section III-A, we summarize the key aspects of this theory as it applies to TCLs, and refer the reader to [12], [13] for full details.

The control objective of [12], [13] is to maintain a TCL group’s on-count between an upper bound and lower bound, while also enforcing each TCL’s individual temperature constraints $\underline{\theta}^i \leq \theta^i \leq \bar{\theta}^i$. The underlying temperature dynamics are the same as (2), except the outdoor temperature θ_a is approximated as constant. Equation (2) is simplified to

$$\frac{d}{dt}\theta^i(t) = \begin{cases} f_{on}^i(\theta^i(t)) & \text{if } \sigma^i(t) = \text{on}, \\ f_{off}^i(\theta^i(t)) & \text{if } \sigma^i(t) = \text{off}. \end{cases} \quad (8)$$

Formally, the on-count of a group of N TCLs is

$$K(t) = \sum_{i=1}^N \mathbb{1}_{\{\text{on}\}}(\sigma^i(t)), \quad (9)$$

and is controlled between the lower bound \underline{K} and the upper bound \bar{K} :

$$\underline{K} \leq K(t) \leq \bar{K}. \quad (10)$$

Control is implemented as an aggregate switching policy $\{\sigma^i\}_{\forall i}$ that determines when a TCL’s power mode should be switched. The switching policy uses time rather than temperature to determine which TCL to switch next. The policy is described in Algorithm 1, where $T_{off}^i(\theta^i(t))$ is the “time to off-exit”, or the time it takes for the i th TCL to travel in the off-mode from its current temperature $\theta^i(t)$ to its upper limit $\bar{\theta}^i$, and $T_{on}^i(\theta^i(t))$ is the time to on-exit, or the time for the i th TCL to travel in the on-mode from $\theta^i(t)$ to its lower limit $\underline{\theta}^i$. We refer to this control strategy as the “Original Strategy”.

Computing a TCL’s time to exit requires data on the TCL’s parameters and real-time measurements of its temperature and power mode. TCL parameters could be communicated off-line, but the real-time communication requirements of the strategy are still high. Methods could be developed to reduce

Algorithm 1 Original Strategy from [12]

- 1: **if** $T_{off}^i(\theta^i(t)) = 0$ **and** TCL i is *off* **then**
 - 2: switch TCL i *on*
 - 3: **end if**
 - 4: **if** $T_{on}^i(\theta^i(t)) = 0$ **and** TCL i is *on* **then**
 - 5: switch TCL i *off*
 - 6: **end if**
 - 7: **if** $K(t^+) < \underline{K}$ for $t^+ > t$ **then**
 - 8: **repeat**
 - 9: switch *on* the TCL in off-mode with the largest time to on-exit
 - 10: **until** $K(t^+) = \underline{K}$ for $t^+ > t$
 - 11: **end if**
 - 12: **if** $K(t^+) > \bar{K}$ **then**
 - 13: **repeat**
 - 14: switch *off* the TCL in on-mode with the largest time to off-exit
 - 15: **until** $K(t^+) = \bar{K}$
 - 16: **end if**
-

the communication burden but there would be a tradeoff in control accuracy.

In [12], the authors derive conditions for values of \underline{K} and \bar{K} that can be satisfied indefinitely by the proposed control strategy. These conditions are useful because they enable us to choose values for \bar{K} and \underline{K} that are known a priori to be feasible. For a group of N TCLs, the condition for lower bound values that can be satisfied indefinitely is

$$\underline{K} < \sum_{i=1}^N \frac{-f_{off}^i(\underline{\theta}^i)/f_{on}^i(\underline{\theta}^i)}{1 - f_{off}^i(\underline{\theta}^i)/f_{on}^i(\underline{\theta}^i)}, \quad (11)$$

and the condition for upper bound values is

$$\bar{K} > N - \sum_{i=1}^N \frac{-f_{on}^i(\bar{\theta}^i)/f_{off}^i(\bar{\theta}^i)}{1 - f_{on}^i(\bar{\theta}^i)/f_{off}^i(\bar{\theta}^i)}. \quad (12)$$

Moreover, it is shown that no other control strategy can do better. That is, no control strategy can achieve a greater lower bound or a smaller upper bound and guarantee that the bound will be satisfied indefinitely. Note that the greatest lower bound (GLB) and the least upper bound (LUB) are equal to the right hand sides of (11) and (12), respectively.

B. Application to Voltage Management Using ACs

1) *Constraining an AC Group’s Power*: For a group of heterogeneous ACs, a count bound can be satisfied by different sets of ACs within the group, and thus can result in different aggregate power levels. We define $\underline{P}_{agg}(\underline{K})$ to be the minimum aggregate power that can occur while satisfying a given lower count bound, and $\bar{P}_{agg}(\bar{K})$ to be the maximum aggregate power that can occur while satisfying a given upper count bound. We find an expression for these power bounds by first ordering the AC group by thermal power rating from least to greatest with index f , i.e., $P_{\Theta}^f \leq P_{\Theta}^{f+1}$. Given a group

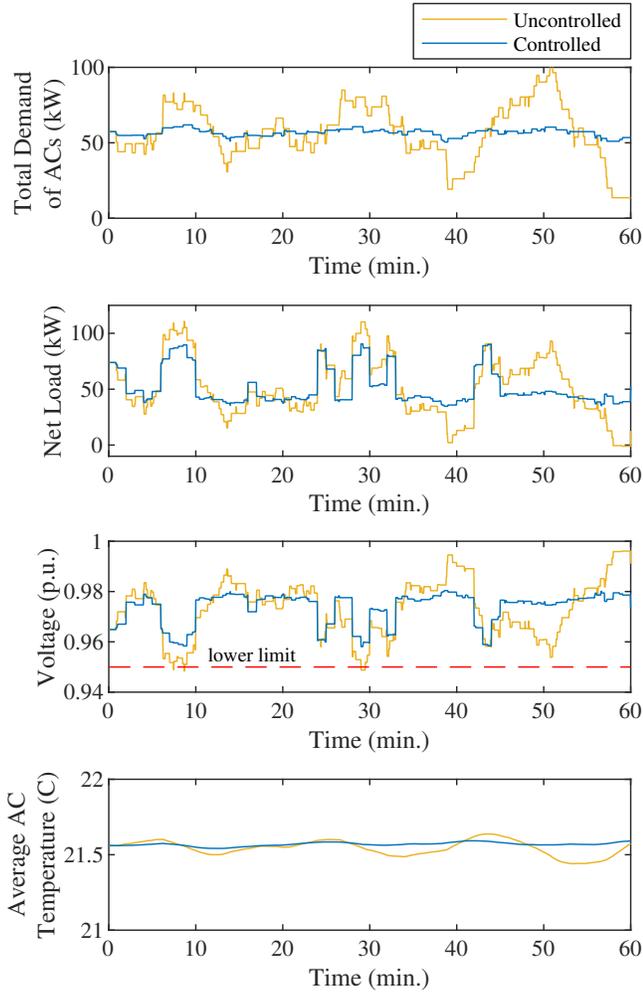


Fig. 5. Comparison of loads and voltage with and without control. The control strategy reduces the variation in the aggregate power consumption of ACs, which reduces net load variation and results in a reduction in voltage excursions. The average temperature of the AC group stays relatively constant in both cases.

of size N and count bounds \bar{K} and \underline{K} , the group's aggregate power will be bounded between

$$P_{\text{agg}}(\underline{K}) = \frac{1}{C_{\text{cop}}} \sum_{f=1}^{\underline{K}} P_{\Theta}^f \quad (13)$$

$$\bar{P}_{\text{agg}}(\bar{K}) = \frac{1}{C_{\text{cop}}} \sum_{f=N-\bar{K}+1}^N P_{\Theta}^f. \quad (14)$$

Note that this computation requires information on the power rating of each AC.

Because our goal is to limit the aggregate variation in power of a group of ACs, we propose selecting \underline{K} and \bar{K} such that the on-count is maximally constrained. In other words, \underline{K} and \bar{K} should be selected such that $(\bar{K} - \underline{K})$ is minimized and $(\bar{K} - \underline{K}) \geq 0$. For larger group sizes, the GLB can be higher than the LUB, in which case the on-count is maximally constrained by setting $\underline{K} = \bar{K}$.

2) *Distribution Simulation with Original Strategy:* Through simulation, we demonstrate the control strategy's

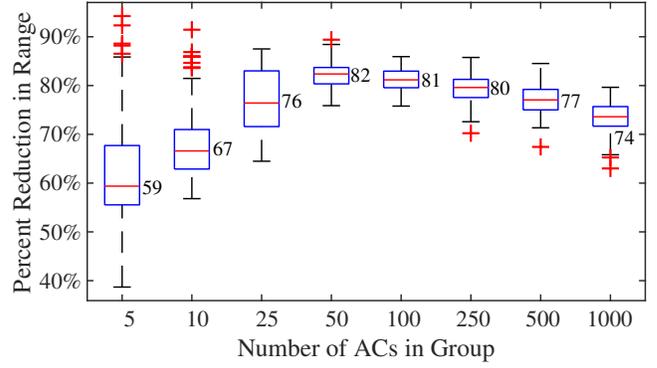


Fig. 6. Reduction in range of power when the Original Strategy is used, across different size groups of ACs. Each box shows the statistics (25th percentile, median, 75th percentile) for 100 Monte Carlo runs at each group size. Medians are also shown numerically. Red + signs are outliers.

ability to constrain the aggregate power of an AC group and thereby reduce voltage excursions on a network. We consider two scenarios: the “Uncontrolled Scenario” in which there are no external switching commands, and the “Controlled Scenario” in which the Original Strategy is used. The simulations use the same models, parameters, and weather data as in Section II-D, but have a duration of only 1 hour. The simulations start at 14:00, which is the hour with the largest range in net load (see Fig. 2). In the Controlled Scenario, the upper and lower count bounds are set such that $\bar{K} = 10$ and $\underline{K} = 10$. Initial temperatures are drawn from the uniform distribution between an AC's upper and lower temperature limit, and the first 10 ACs are initialized to the on mode.

Results for the two scenarios are compared in Fig. 5: the total demand of the AC group is much less variable in the Controlled Scenario than the Uncontrolled Scenario (top plot), which results in less variability in net load (second plot), which in turn results in less variability in voltage (third plot). The bottom plot shows that variation in the AC group's average temperature is very small, both with and without control.

3) *Monte Carlo Simulation with Original Strategy:* Finally, we use Monte Carlo methods to explore the effectiveness of the control strategy at reducing aggregate power variation for different size groups of ACs. We run 100 simulations for each of eight group sizes ranging from 5 to 1000 ACs. For each simulation, AC parameters and initial conditions are re-sampled from their respective random distributions. For all simulations, the simulation duration is 12 hours, the outdoor temperature is 32°C, and one third of ACs are initialized in the on-mode. We simulate the ACs for an hour prior to the test to ensure steady-state conditions have been reached.

Figure 6 shows that, for each group size, the control strategy reduces the range of aggregate power. Each box plot represents a distribution across 100 Monte Carlo simulations. For group sizes under 50, the median percent reduction in range decreases with decreasing group size because maintaining $\underline{K} = \bar{K}$ is infeasible for some smaller-sized groups.

Instead, for these smaller groups, the on-count is maintained between two distinct bounds, for example, $\bar{K} = 4$ and $\underline{K} = 3$ for a group of 10. For group sizes of 50 or more, the range in the on-count decreases by 100% because setting $\underline{K} = \bar{K}$ is feasible for all groups of this size. However, the range in power cannot decrease by 100% because of TCLs' heterogeneous power ratings (see (13)-(14)). For group sizes greater than 50, the controlled range in power grows with increasing group size, which results in a slow decline in the percent reduction in range. Thus, we find that the control strategy is most effective at reducing the range in power of medium sized groups, but, for all group sizes tested, one can expect at least a 40% reduction in range.

IV. CONTROL STRATEGY FOR TCLs WITH LOCKOUT

A drawback of the Original Strategy is the possibility of switching a TCL too frequently, which could damage the TCL's compressor [23]. In theory, the strategy can switch a TCL arbitrarily frequently; in practice, it is limited by the simulation time step. For example, in the Monte Carlo simulations of groups of 5 ACs the minimum time between switching for an individual AC was 2 sec., which was the simulation time step. Many TCLs have a built-in "lockout period" that enforces a minimum time period between switching. Because the Original Strategy does not account for lockout constraints, it can fail when applied to a group of TCLs with lockout. Thus, we propose a new switching policy for TCLs that have lockout constraints. Lockout constraints are a current thrust of research on TCL control, and our proposed control strategy adds to a growing body of work including [24]–[27].

A. Lockout Modeling

We include lockout constraints within the TCL model by adding a state that indicates when a TCL is locked or unlocked. Given a lockout period of t_L for all TCLs, the lockout state λ^i for the i th TCL is determined by

$$\lambda^i(t) = \begin{cases} \text{locked} & \text{if } \int_{t-t_L}^t \mathbb{1}_{\{\sigma^i(\tau)\}}(\sigma^i(\tau)) d\tau < t_L, \\ \text{unlocked} & \text{otherwise.} \end{cases} \quad (15)$$

If a TCL is locked, then its power mode cannot be switched. This condition is enforced by

$$\sigma^i(t) = \sigma^i(t^-) \quad \text{if } \lambda^i(t) = \text{locked}. \quad (16)$$

B. Control Strategy for Lockout

We propose a new control strategy for a group of TCLs with lockout that bounds the group's on-count while also satisfying individual temperature constraints. We design the strategy to avoid system states that would lead to constraint violations. For example, to prevent more than \bar{K} TCLs from being locked *on* at the same time, no more than \bar{K} TCLs can be in the *off* mode with $T_{\text{off}}^i(\theta^i(t)) < t_L$. The proposed strategy prevents this problematic system state by switching TCLs well before they reach their upper temperature limits. Specifically, the strategy begins to try to switch a TCL once

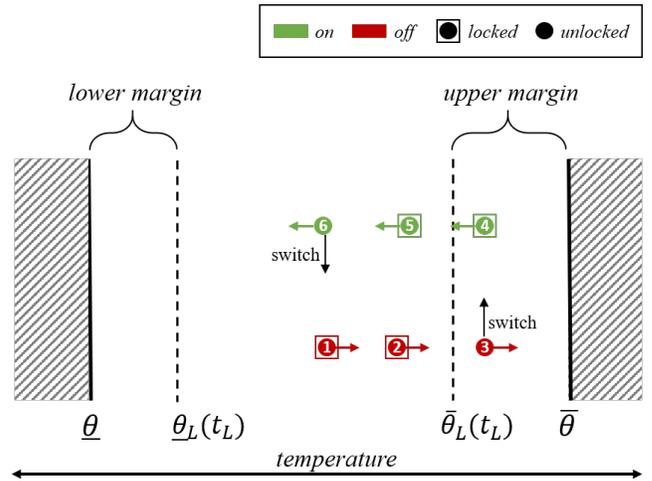


Fig. 7. Illustration of an AC's temperature range, on and off modes, locked status, and the upper and lower margins used in the control strategy. The numbering 1-6 indicates different snapshots in time to illustrate the AC progressing through its state space when controlled by the Lockout Strategy.

it has entered the "upper temperature margin" or "lower temperature margin" (see Fig. 7).

We formally describe the control strategy in Algorithm 2 and refer to it as the "Lockout Strategy". The strategy uses the following definitions:

- $\bar{\theta}_L^i$ is the inner temperature of the upper temperature margin for TCL i , and is explicitly defined below;
- $\underline{\theta}_L^i$ is the inner temperature of the lower temperature margin for TCL i , and is explicitly defined below;
- $\bar{M}(t)$ is the set of *off* TCLs in the upper margin;
- $\underline{M}(t)$ is the set of *on* TCLs in the lower margin;
- $A_{\text{on} \rightarrow \text{off}}(t)$ is the set of TCLs available to switch *off*;
- $A_{\text{off} \rightarrow \text{on}}(t)$ is the set of TCLs available to switch *on*.

A TCL's temperature margins are a function of the TCL's thermal parameters and lockout time. For the i th TCL, $\underline{\theta}_L^i = \max(\theta_{L1}^i, \theta_{L2}^i)$, where θ_{L1}^i is the temperature reached after the TCL travels in the off-mode from $\underline{\theta}^i$ for a t_L length of time and θ_{L2}^i is the starting temperature from which it takes a t_L length of time for the TCL to travel in the on-mode to $\underline{\theta}^i$. For the upper margin, $\bar{\theta}_L^i = \min(\theta_{L3}^i, \theta_{L4}^i)$, where θ_{L3}^i and θ_{L4}^i are defined similarly but with respect to $\bar{\theta}^i$.

We define "available to switch" in a way that protects lockout constraints. For example, an *on* TCL should not be available to switch if it is too close to $\bar{\theta}^i$ because, once in the *off* mode, it may take less than the lockout period to travel back to $\bar{\theta}^i$. A similar argument can be made for *off* TCLs with respect to $\underline{\theta}^i$. We ensure lockout constraints are satisfied by defining a TCL to be available to switch if it is *on*, *unlocked*, and $\theta^i(t) < \bar{\theta}_L^i$ or it is *off*, *unlocked*, and $\theta^i(t) > \underline{\theta}_L^i$.

By comparing Algorithms 1 and 2, one can see that the Lockout Strategy differs from the Original Strategy in a number of ways. The Lockout Strategy not only has awareness of which TCLs are locked but also has features to avoid the undesirable effects of lockout. The most significant

of these features is switching TCLs at the beginning of their temperature margins instead of at their temperature limits. Another feature is the use and definition of “available to switch”. Yet another feature is reserving “switching capacity” for TCLs that are locked and in sets $\underline{M}(t)$ or $\overline{M}(t)$ so that, once *unlocked*, they can switch right away. We reserve switching capacity by ordering TCLs in the margin by their time to exit (see lines 2 and 10 of Algorithm 2) and only switching a TCL with locked TCLs ahead of it (i.e., with a shorter time to exit) if there is enough space between the bound and $K(t)$ or there are enough TCLs available to switch in the opposite direction (see lines 3, 5, 11, 13).

Algorithm 2 Lockout Strategy

- 1: **if** $|\overline{M}(t)| > 0$ **then**
 - 2: order TCLs in \overline{M} by time to *off*-exit in ascending order, and set m equal to the smallest index in \overline{M} for which the TCL is also *unlocked*
 - 3: **if** $m \leq \overline{K} - K(t) + |A_{\text{on} \rightarrow \text{off}}(t)|$ **and** $\overline{K} - K(t) > 0$ **then**
 - 4: switch TCL m *on*
 - 5: **else if** $m \leq \overline{K} - K(t) + |A_{\text{on} \rightarrow \text{off}}(t)|$ **then**
 - 6: switch TCL m *on* and simultaneously switch *off* the TCL in $A_{\text{on} \rightarrow \text{off}}(t)$ with the largest time to *off*-exit
 - 7: **end if**
 - 8: **end if**
 - 9: **if** $|\underline{M}(t)| > 0$ **then**
 - 10: order TCLs in \underline{M} by time to *on*-exit in ascending order and set m equal to the smallest index in \underline{M} for which the TCL is also *unlocked*
 - 11: **if** $m \leq K(t) - \underline{K} + |A_{\text{off} \rightarrow \text{on}}(t)|$ **and** $K(t) - \underline{K} > 0$ **then**
 - 12: switch TCL m *off*
 - 13: **else if** $m \leq K(t) - \underline{K} + |A_{\text{off} \rightarrow \text{on}}(t)|$ **then**
 - 14: switch TCL m *off* and simultaneously switch *on* the TCL in $A_{\text{off} \rightarrow \text{on}}(t)$ with the largest time to *on*-exit
 - 15: **end if**
 - 16: **end if**
 - 17: **if** $T_{\text{off}}^i(\theta^i(t)) = 0$ **and** TCL i is *off* **then**
 - 18: switch TCL i *on*
 - 19: **end if**
 - 20: **if** $T_{\text{on}}^i(\theta^i(t)) = 0$ **and** TCL i is *on* **then**
 - 21: switch TCL i *off*
 - 22: **end if**
-

C. Constraints on Initial Conditions

Prior to the start of the strategy, we assume that TCLs have been operating “naturally”, i.e., without external control. This implies that an *off* TCL will only be locked if $\theta^i(t_0) < \theta_{L1}^i$ and an *on* TCL will only be locked if $\theta^i(t_0) > \theta_{L3}^i$, where t_0 is the time at which the strategy starts.

For the control strategy to succeed, TCLs’ initial states must satisfy the following conditions. First is a trivial condition: TCLs’ initial temperatures and the group’s on-count must satisfy their constraints. Second, to prevent more than \overline{K} from being locked *on* at the same time, no more than \overline{K}

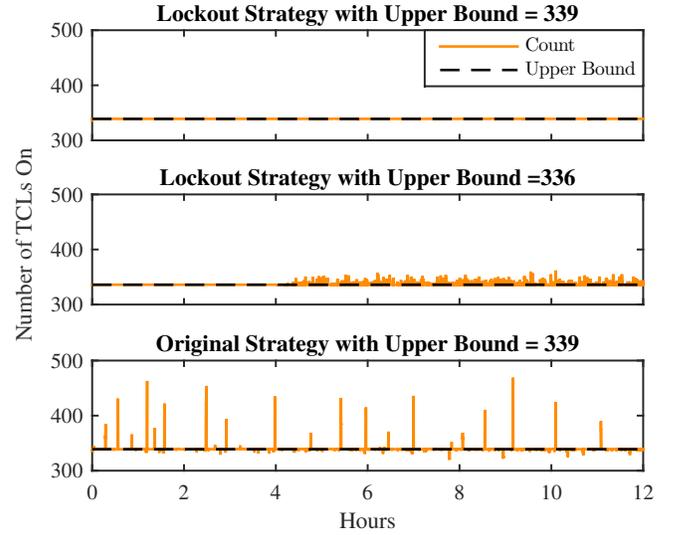


Fig. 8. Simulations of a group of 1000 ACs with lockout, controlled to satisfy an upper count bound. The Lockout Strategy is able to satisfy the upper bound of 339 computed with (18) (top) but unable to satisfy the upper bound of 336 computed with (12) (middle). When we use the Original Strategy and the upper bound computed with (18), the strategy is unable to satisfy the upper bound (bottom).

TCLs should be in the *off* mode with $T_{\text{off}}^i(\theta^i(t)) < t_L$. Third, to prevent fewer than \underline{K} from being *on*, no more than $(N - \underline{K})$ TCLs should be in the *on* mode with $T_{\text{on}}^i(\theta^i(t)) < t_L$.

D. Counting Bounds Adjusted for Lockout

We hypothesize that, under the Lockout Strategy, the on-count bounds that can be satisfied indefinitely will satisfy conditions similar to those in (11)-(12). Instead of depending on the upper and lower temperature limits, we expect the conditions will depend on the margin temperatures $\underline{\theta}_L$ and $\overline{\theta}_L$, as formulated below.

$$\underline{K} < \sum_{i=1}^N \frac{-f_{\text{off}}^i(\underline{\theta}_L^i)/f_{\text{on}}^i(\underline{\theta}_L^i)}{1 - f_{\text{off}}^i(\underline{\theta}_L^i)/f_{\text{on}}^i(\underline{\theta}_L^i)} \quad (17)$$

$$\overline{K} > N - \sum_{i=1}^N \frac{-f_{\text{on}}^i(\overline{\theta}_L^i)/f_{\text{off}}^i(\overline{\theta}_L^i)}{1 - f_{\text{on}}^i(\overline{\theta}_L^i)/f_{\text{off}}^i(\overline{\theta}_L^i)} \quad (18)$$

We are currently working on a proof of correctness for the Lockout Strategy. We aim to show that the strategy will satisfy the bounds \overline{K} and \underline{K} indefinitely if the bounds satisfy conditions (17)-(18).

E. Simulation Results

We demonstrate the Lockout Strategy with simulations of a group of 1000 heterogeneous ACs with lockout; for now, we omit a distribution system model but will include it in future work. AC parameters are randomly selected as described in Section II-B, and each AC is given a 1 minute lockout time. We run three simulations: in the first, the Lockout Strategy is tested with \overline{K} set to the “adjusted” LUB determined by (18); in the second, the Lockout Strategy is tested with \overline{K} set to the original LUB determined by (12); and in the third,

the Original Strategy is tested with the adjusted LUB. In all simulations, the AC group is initialized such that $\theta^i(t_0) = \bar{\theta}_L^i$ for all TCLs and $K(t_0) = \bar{K}$. The lower bound \underline{K} is set to zero for all simulations.

As shown in the top plot of Fig. 8, the Lockout Strategy is able to satisfy the adjusted LUB for the simulated 12 hours. However, when the upper bound decreases to the original LUB, the strategy is unable to satisfy the bound (see middle plot). Finally, the Original Strategy is unable to satisfy the adjusted LUB (see bottom plot).

Voltage management may suffer if a non-lockout strategy is used for a group of TCLs that has lockout constraints. As the bottom plot of Fig. 8 shows, the Original Strategy, which does not account for lockout, causes large spikes in a group's on-count. Although the simulation does not include power flow, on-count deviations from the bound indicate aggregate power deviations, which could cause voltage excursions on an actual distribution feeder.

V. CONCLUSIONS

We have demonstrated a novel TCL control strategy and its application to distribution voltage management. The strategy minimizes variation in the number of TCLs that are on by constraining the on-count between close lower and upper bounds. We have shown that the strategy is able to reduce the range of TCLs' aggregate power, and thereby reduce voltage excursions on a distribution system. We found that the control strategy is most effective at reducing the range in aggregate power of medium-size groups of TCLs (i.e., 50-100 TCLs).

We have also proposed a new control strategy for TCLs with lockout and have demonstrated the strategy's ability to satisfy on-count bounds over 12 consecutive simulation hours. In future work, we will investigate heterogeneous lockout periods and the effect of longer lockout periods on the set of feasible bounds. In addition, we plan to incorporate the proposed strategy into a broader control architecture to track a balancing signal with aggregated TCLs while still ensuring that local constraints, such as voltage, are satisfied.

REFERENCES

- [1] E. C. Kern and E. M. Gulachenski, "Cloud effects on distributed photovoltaic generation: Slow transients at the gardner, massachusetts photovoltaic experiment," *IEEE Trans. on Energy Conversion*, vol. 4, no. 2, pp. 184–190, 1989.
- [2] W. Kersting, *Distribution System Modeling and Analysis*, 3rd ed. Boca Raton, Florida: CRC Press, 2012.
- [3] N. A. Sinityn, S. Kundu, and S. Backhaus, "Safe protocols for generating power pulses with heterogeneous populations of thermostatically controlled loads," *Energy Conversion and Management*, vol. 67, pp. 297–308, 2013.
- [4] T. Voice, "Stochastic thermal load management," *IEEE Trans. on Automatic Control*, vol. 63, no. 4, pp. 931–946, 2017.
- [5] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proc. of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan. 2011.
- [6] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Trans. on Power Systems*, vol. 28, no. 1, pp. 430–440, Feb. 2013.
- [7] S. Bashash and H. Fathy, "Modeling and control of aggregate air conditioning loads for robust renewable power management," *IEEE Trans. on Control Syst. Tech.*, vol. 21, no. 4, pp. 1318–1327, 2013.

- [8] W. Zhang, J. Lian, C. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Trans. on Power Systems*, vol. 28, no. 4, pp. 4655–4664, Nov. 2013.
- [9] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *IEEE International Conf. on Smart Grid Comm.*, Oct. 2013.
- [10] M. Chertkov, D. Deka, and Y. Dvorkin, "Optimal ensemble control of loads in distribution grids with network constraints," 2017. [Online]. Available: <http://arxiv.org/abs/1710.09924>
- [11] E. Dall'Anese, S. S. Guggilam, A. Simonetto, Y. C. Chen, and S. V. Dhople, "Optimal regulation of virtual power plants," *IEEE Trans. on Power Systems*, vol. 33, no. 2, pp. 1868–1881, Mar. 2018.
- [12] P. Nilsson and N. Ozay, "On a class of maximal invariance inducing control strategies for large collections of switched systems," in *Proc. of the International Conference on Hybrid Systems: Computation and Control (HSCC)*, Apr. 2017.
- [13] —, "Maximizing the time of invariance for large collections of switched systems," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, Dec. 2017.
- [14] A. P. Dobos, "PVWatts version 5 manual," National Renewable Energy Laboratory, 2014.
- [15] R. Fu, D. Feldman, M. Woodhouse, and K. Ardani, "U.S. solar photovoltaic system cost benchmark: Q1 2017," National Renewable Energy Laboratory, Tech. Rep. NREL/TP-6A20-68925, 2017.
- [16] C. Maxey and A. Andreas, "Rotating shadowband radiometer (RSR)," Oak Ridge National Laboratory, Tech. Rep. NREL DA-5500-56512, 2007.
- [17] R. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Ph.D. dissertation, Princeton Univ., N.J., 1978.
- [18] D. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, pp. 1389–1400, 2009.
- [19] P. Du and N. Lu, "Appliance commitment for household load scheduling," *IEEE Trans. on Smart Grid*, vol. 2, no. 2, pp. 411–419, 2011.
- [20] M. Cohen and D. Callaway, "Effects of distributed PV generation on California's distribution system, Part 1: Engineering simulations," *Solar Energy*, vol. 128, pp. 126–138, 2016.
- [21] GridLAB-D software. [Online]. Available: <http://www.gridlabd.org/>
- [22] National Electrical Manufacturers Association, "American national standard for electric power systems and equipment – voltage ratings (60 Hertz)," Tech. Rep. ANSI C84.1-2006, Dec. 2006.
- [23] B. H. Pinckaers, "Minimum off-time circuit," Honeywell, US Patent 3,619,668, Nov. 1971.
- [24] L. C. Totu, R. Wisniewski, and J. Leth, "Demand response of a TCL population using switch-rate actuation," *IEEE Trans. on Control Systems Technology*, vol. 25, no. 5, pp. 1537–1551, 2017.
- [25] 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)*, Jun. 2018.
- [26] C. Y. Chang, W. Zhang, J. Lian, and K. Kalsi, "Modeling and control of aggregated air conditioning loads under realistic conditions," *IEEE PES Innovative Smart Grid Technologies Conference*, 2013.
- [27] L. A. D. Espinosa, M. Almassalkhi, P. Hines, S. Heydari, and J. Frolik, "Towards a macromodel for packetized energy management of resistive water heaters," in *Conference on Information Sciences and Systems (CISS)*, Mar. 2017.