

Heuristic Algorithms for Aggregated HVAC Control via Smart Thermostats for Regulation Service

Rajendra Adhikari¹, *Member, IEEE*, Manisa Pipattanasomporn², *Senior Member, IEEE*,
and Saifur Rahman³, *Life Fellow, IEEE*

Abstract—Residential HVAC control is a large untapped resource for providing regulation services to the grid. This paper presents a set of algorithms for controlling HVACs of a group of residential houses that a demand response aggregator can use to sell regulation service in the wholesale market. The focus is on the regulation market offered by the PJM RTO. Real-world regulation signals from PJM are used to simulate the performance and range of regulation services in a realistic scenario. After presenting the empirical counter example for why a universal optimal control strategy cannot exist for regulation, a set of heuristic algorithms is presented, which performs well in a range of test cases. The control mechanism involves a central controller communicating with smart thermostats of multiple residential houses to gather indoor temperature data, prioritizing them according to certain heuristics and sending on/off signals back to the thermostats to control the HVAC. The case studies indicate that the proposed heuristic algorithms can deliver the required regulation services, while adequately handling communication delays, different types of regulation signals and household's thermal comfort requirements.

Index Terms—Regulation services, aggregated HVAC control, smart thermostat.

I. INTRODUCTION

REGIONAL transmission operators (RTO), such as PJM in the U.S., are responsible for managing the wholesale electricity market for smooth operation of the electric grid. In PJM, the wholesale markets consist of: (i) the energy market—which ensures real-time balance between the energy produced and consumed; (ii) the capacity market—which ensures sufficient generation capacity to serve the expected system load; and (iii) the ancillary services market—which provides the reserve and regulation services [1]. A regulation service requires participants to vary their generation (or consumption in the case of a demand response (DR) resource) in response to regulation signals dispatched by RTO. Two kinds of regulation signals are available in the PJM market: slow changing traditional regulation signal (RegA) and fast

changing (every 2 seconds) dynamic regulation signal (RegD). Participants can bid the amount, price and type of regulation services (RegA or RegD) for each hour of the next day. And, if they win the bid, they are expected to adjust their generation (or consumption) to closely follow the regulation signals during that hour.

One type of participants for a regulation service is a DR aggregator who aggregates at least a certain amount of building-level loads (e.g., at least 0.1MW for PJM) and control them to make the aggregated power follow the regulation signal. Typical building-level loads that can be aggregated/controlled are a group of residential Heating, Ventilation and Air Conditioning (HVAC) units. Aggregated control of HVACs has a number of unique challenges. First challenge comes from the requirement to keep the indoor temperature of each house within an acceptable homeowner's comfort range. This constraint forces HVACs to turn ON or OFF to maintain the indoor temperature around the pre-set set point (e.g., 77°F). And, with simultaneous operation of many HVACs, it will create inadvertent and large aggregated power fluctuation that interferes with the objective of controlling their aggregated power consumption. The second challenge comes from the fact that AC compressors cannot be cycled very quickly due to their built-in protection mechanism to ensure the minimum amount of time an HVAC is turned ON. This poses a challenge in controlling the aggregated power of HVACs to track fast moving regulation signals.

To tackle these challenges, methods to control a group of HVACs were discussed in the literature. In [2], researchers presented a mechanism (called safe protocol) to produce positive and negative power pulses without disturbing the state diversity and creating peaks due to inadvertent synchronization. In [3], the authors proposed a temperature priority-based algorithm to control an aggregation of HVAC units for load following where compressor lockouts (because of the minimum ON/OFF time constraint) are taken into account. The work was extended in [4] to take into account thermal behavior variations among different houses. The work however was done using a one-minute signal as opposed to the two-second regulation signal. Also, the HVAC rated capacity was assumed to be the same for all the houses. The weighted minimization problem of the differences between wind power and residential demand and between the desired temperature and the actual temperature was explored in [5]. However, the model assumed continuous variability of HVAC power. It also ignored minimum ON/OFF time requirements or problem of inadvertent

Manuscript received September 19, 2018; revised January 10, 2019, May 3, 2019, and July 29, 2019; accepted August 31, 2019. Date of publication October 8, 2019; date of current version April 21, 2020. Paper no. TSG-01373-2018. (*Corresponding author: Rajendra Adhikari.*)

R. Adhikari and S. Rahman are with the Advanced Research Institute, Bradley Department of Electrical and Computer Engineering, Virginia Tech, Arlington, VA 22203 USA (e-mail: rajee@vt.edu).

M. Pipattanasomporn is with the Smart Grid Research Unit, Department of Electrical Engineering, Chulalongkorn University, Bangkok 10330, Thailand.

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

Digital Object Identifier 10.1109/TSG.2019.2946170

1949-3053 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

synchronizations. Authors in [6] have added frequency control into a simple dead-band based HVAC control as a means to provide frequency response to the Great Britain grid. The work disregarded hard temperature limits and the aggregation effect was not properly tackled to make the power follow a regulation signal. Authors in [7] presented a state-space model for aggregated HVAC control that could be used for load following. However, it was a simplified aggregated model that did not consider individual HVAC dynamics and minimum ON/OFF time requirements. In [8], the sliding mode set point change-based controller was used for making the aggregated HVAC power follow Automatic Generation Control (AGC) signals. The meta-heuristic natural aggregation algorithm was used to solve the HVAC status scheduling problem for load-shape modification and cost-minimization in [9]. But the approach was only applicable for cases when the aggregated HVAC power needed to follow a known power profile, which is not the case for constantly changing regulation signals.

Some work was carried out in the area of commercial HVAC control. Regulation services were explored in the case with variable and controllable fans [10], [11]. A decentralized control algorithm for HVACs for load-shape management was explored in [12] and a partial differential equation-based aggregated model incorporating set-point change was developed in [13]. However, these aggregated mathematical models did not consider ON/OFF time requirements of individual compressors. While commercial HVACs were clearly identified as a great resource for regulation services [14], the residential sector on the other hand provides untapped demand-side resources and is expected to contribute significantly in the regulation market of the future [15].

The work most closest to ours is found in [16] where the AC compressor's minimum ON/OFF times were explicitly modelled in the state-space model. Authors tackle unpredictability of regulation signals by maximizing the 'regulation capacity', defined to be proportional to the number of unlocked HVACs. The state queuing model employed in this and other similar work [17], [18] to make control decision involved several modelling simplifications that made even enormous problem size tractable; however, this paper employed the HVAC physical model instead of the queuing model for better accuracy [19]. While authors in [20] presented an aggregated HVAC control mechanism similar to ours, but we use time-to-boundary based approach compared to temperature deviation based approach used by the authors.

In this paper, we tackle the aggregated HVAC control problem from the first principle using the second order thermal dynamics, and present an intuitive empirical proof that when the minimum compressor ON/OFF time constraint is imposed, no algorithm can perform optimal control for all kinds of regulation signals. This provides theoretical clarity to the problem and sets up a realistic expectation for what can and cannot be achieved by using algorithms and optimization methods. We present a set of intuitive heuristic algorithms that is demonstrated to perform well in tracking fast moving (two-second intervals) real-world RegD signals from PJM. A method to determine the regulation capability to bid to the wholesale market to maximize the profit is also presented. In addition,

the impact of different comfort settings, communication delay and the regulation signal type on the regulation performance score and credit has also been analyzed.

Hence, the contributions of this paper include:

- i) Empirical proof by counter-example for why universal optimal regulation algorithm cannot exist.
- ii) A complete framework for implementing a residential HVAC control based regulation service for DR aggregators, including the process of determining the regulation capability.
- iii) A set of log-linear (on the number of HVACs) real-time algorithm for controlling the group of HVACs to closely follow the regulation signal that has good performance on wide range of real-world regulation signals.
- iv) Analysis of the impact of communication delay, signal type and homeowner comfort preference on the regulation service based on the performance score calculation method defined by PJM.

II. FRAMEWORK AND PROBLEM FORMULATION

In this study, a DR framework similar to that proposed in [21] was used where regulation services were performed by a DR aggregator by controlling aggregated household loads. Each house was assumed to have one HVAC unit. In this framework, a central controller (belonging to the DR aggregator) gathers information from smart thermostats required to make decisions about which HVAC to control using algorithms presented in Section IV. The controller then sends the signal back to the thermostats to turn ON/OFF the HVACs. The objective here is to come up with a set of methods and algorithms for the aggregator to control the collection of HVACs so as to gain maximum credit.

A. Calculation of Performance Score (PS) and Regulation Service Pay-Off

PJM allows for using sub-metered data to verify the delivery of regulation services [22]. Hence, it is sufficient to collect HVAC power consumption data from each participating house and submit their aggregation as a proof of regulation service delivery. DR resources must provide at least 0.1 MW of regulation capability (RegMW) and should submit their midpoint MW value [22], [23].

Fig. 1 illustrates the basic principle of regulation service provision. Aggregated load must be controlled, as closely as possible, to follow the two-second regulation signals. The aggregated load varies around a value, called MidpointPower, through a certain range, called the regulation capability (RegMW). PJM ranks and pays regulation resources based on the performance score (PS) and RegMW. The more closely and quickly the aggregated power follows the regulation signal, the higher is the PS. PJM provides mathematical formulae to calculate PS in [23].

In PJM, the pay-off for regulation services, i.e., Total Regulation Credit (TRC), is determined as the sum of Capacity Credit ($C_{capacity}$) and Performance Credit ($C_{performance}$) [23], where:

$$C_{capacity} = RegMW \times PS \times CCP \quad (1)$$

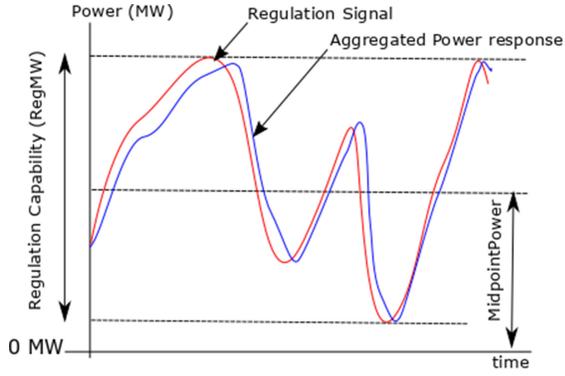


Fig. 1. Illustration of regulation service.

$$C_{performance} = RegMW \times PS \times MR \times PCP \quad (2)$$

$$TRC = C_{capacity} + C_{performance} \quad (3)$$

The CCP (Capability Clearing Price) and PCP (Performance Clearing Price) are determined by PJM through the wholesale market clearing process and can be considered independent variables for the optimization purpose. MR (Hourly Mileage Ratio) is the ratio of the hourly mileage of the regulation signal to the 30-day average of mileage of RegA and is an independent variable calculated based on historical regulation signal [24].

B. Maximizing Regulation Service Pay-Off

The objective is to maximize the TRC. Assuming that the regulation service is to be provided for a period T (from t_{start} to t_{end}). To maximize the total credit which is proportional to $PS \times RegMW$ (as per (3)), the objective can be stated as:

$$\text{maximize } PS \times RegMW$$

$$\text{Subject to: } \theta_{lower_n} \leq T_{A_n}^{t_k} \leq \theta_{upper_n} \quad \forall n, \forall k$$

$$response_{t_k} = \sum_{n=1}^N P_{HVAC_n} * U_n^{t_k} \quad \forall k$$

$$T_{A_n}^{t_{k+1}} = f(T_{A_n}^{t_k}, T_{M_n}^{t_k}, C^{t_k}, T_o^{t_k}, \Delta t, U_n^{t_k})$$

$$U_n^{t_k} = 1 \text{ if } U_n^{t_{k-1}} = 1 \text{ and } t_k - t_{last_ON}^n \leq MIN_ON$$

$$U_n^{t_k} = 0 \text{ if } U_n^{t_{k-1}} = 0 \text{ and } t_k - t_{last_OFF}^n \leq MIN_OFF$$

(4)

where,

$RegMW$	The regulation capability provided (kW)
PS	The performance score, calculated based on response and regulation signal as per [23].
t_k	Time step. Duration T is divided into a series of time steps $t_{start} \leq t_k < t_{end}$
$T_{A_n}^{t_k}$	Indoor air temperature of house n at time t_k ($^{\circ}F$)
$\theta_{lower_n}, \theta_{upper_n}$	Lower/upper bounds of acceptable temperature of house n ($^{\circ}F$)
P_{HVAC_n}	Rated power of HVAC at house n (kW)
$U_n^{t_k}$	HVAC state (1=ON/0=OFF) for house n at time t_k (and the HVAC is referred to as $HVAC_n$)

$T_{A_n}^{t_{k+1}}$	Air temperature in the next time step-house n ($^{\circ}F$)
f	A function that models second order thermal dynamics of a house and expresses indoor air temperature in the next time step
$T_{A_n}^{t_k}$	Air temperature at time step t_k ($^{\circ}F$)
$T_{M_n}^{t_k}$	Building mass temperature at time step t_k ($^{\circ}F$)
C^{t_k}	House thermal parameters (e.g., insulation, heat gains and thermal capacity) at time step t_k
$T_o^{t_k}$	Outdoor air temperature at time step t_k ($^{\circ}F$)
Δt	The interval between two time steps
$t_{last_ON}^n$	The latest time step before t_k when HVAC n was turned ON
$t_{last_OFF}^n$	The latest time step before t_k when HVAC n was turned OFF
MIN_ON	The minimum time for which an HVAC must run once it is turned on (i.e., two minutes in this study)
MIN_OFF	The minimum time for which an HVAC must remain off once it is turned OFF (i.e., three minutes in this study).

Since PS is a measure of how closely the response signal follows the regulation signal (see Fig. 1), it is dependent on RegMW, MidpointPower and the HVAC state ($U_n^{t_k}$) during each control interval, and these are the decision variables. Adjustment of HVAC states ($U_n^{t_k}$) must satisfy minimum ON/OFF time requirements, and at the same time, ensuring indoor temperature comfort constraints. The objective in plain words is to find the best RegMW and control strategy for all participating HVACs so that the TRC is maximized, while meeting the comfort requirements and device constraints of all houses.

In this study, the second order equivalent thermal parameter (ETP) model, i.e., the two-node (building mass, and building air) heat transfer model from [21], [22], was used to determine function f . HVACs were modelled as a constant power consumption device with a constant heat removal rate (when ON). Their rated capacity was randomly varied among different houses.

III. PROOF FOR NON-EXISTENCE OF OPTIMAL SOLUTION

In our previous work [21], an optimal control strategy for a group of HVACs was derived to keep the aggregated HVAC power below a fixed minimum possible level during a DR period while respecting homeowners' comfort constraints. One might think that a similar optimal strategy might exist that makes the aggregated HVAC power optimally follow a regulation signal while respecting comfort constraints. However, regulation signals change every two seconds, which is much shorter than the minimum HVAC ON/OFF time constraint (e.g., 2-3 minutes) introduced to prevent short cycling of AC compressors. Therefore, no algorithm can preemptively perform optimal control of a group of HVACs to make the aggregated power follow a regulation signal most closely. The proof-by-counterexample for non-existence of optimal solution is presented next.

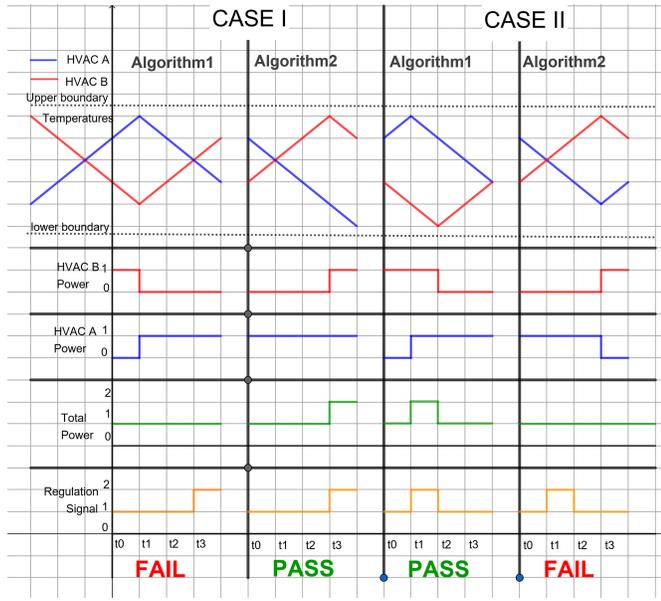


Fig. 2. HVAC control illustration.

Consider a simplified regulation problem with just two HVACs, HVAC A in house A and HVAC B in house B. The rated power of both HVACs was assumed to be one unit and the minimum ON/OFF time was set to be three time steps. The HVACs were considered to be in their cooling mode. At the beginning of the regulation period at t_0 , both HVACs were assumed to have fulfilled their minimum ON/OFF time requirements, and free to change their states. Two different cases of regulation signals were considered, Case I and Case II, as outlined at the bottom of Fig. 2.

First, let us consider Case I.

At t_0 : the regulation signal dictates that the aggregated power be equal to one unit. There can only be two algorithms that can meet this requirement, Algorithm 1 which turns on HVAC B at t_0 and Algorithm 2 which turns on HVAC A at t_0 . No other scenario is possible.

At t_1 : Algorithm 1 turns on HVAC A and turn off HVAC B to prevent house A's indoor temperature from hitting the upper boundary, while Algorithm 2 is forced to keep HVAC A turned on to maintain its minimum ON time requirement.

At t_2 : the states of HVACs need to be maintained to be same as that at t_1 to fulfill the minimum ON/OFF time requirements. Up till t_2 , the aggregated power matches the regulation signal.

At t_3 : the regulation signal increases the required aggregated HVAC power needs to two units. This requires both HVACs to turn ON. Algorithm 2 can do it by keeping HVAC A on and turning on HVAC B as well, which by now has completed minimum off time requirement. Algorithm 1, however, cannot turn on HVAC B as that would violate the minimum OFF time requirement. Thus, for Case I, Algorithm 2 can follow the regulation signal, and is optimal, whereas Algorithm 1 cannot.

Now, let us consider Case II.

At t_0 : the regulation signal is same as before (at one unit), so both algorithms behave exactly as before.

Algorithm 1 Greedy Algorithm (GA) for HVAC Control

```

1: Get  $RegMW$ 
2: for each time step  $k$ :
3:   for each HVAC  $n$ :
4:     Calculate  $B_n^k$  [21]
5:   end for
6:    $sorted\_hvac \leftarrow$  sort list of
   HVACs based on  $B_n^k$ 
7:    $sum = 0$ 
8:    $full = false$ 
9:   for HVAC $_n$  in  $sorted\_hvac$ :
10:    if  $mustRun(HVAC_n)$ :
11:      $U_n^k = 1$ 
12:      $sorted\_hvac.remove(HVAC_n)$ 
13:      $sum = sum + P_{HVAC_n}$ 
14:     if  $mustNotRun(HVAC_n)$ :
15:       $U_n^k = 0$ 
16:      $sorted\_hvac.remove(HVAC_n)$ 
17:   for HVAC $_n$  in  $sorted\_hvac$ :
18:    if  $B_n^k + D_n^k \leq B_n^{max}$  and not full:
19:     if  $P_{HVAC_n} + sum \leq D_L^k$ :
20:       $U_n^k = 1$ 
21:       $sum = sum + P_{HVAC_n}$ 
22:     else:
23:       $U_n^k = 0$ 
24:       $full = true$ 
25:     else:
26:       $U_n^k = 0$ 
27:   end for
28: end for

```

Algorithm 2 Lazy Algorithm (LA) for HVAC Control

Insert the following line in Algorithm 1 between line 6 and 7:
7: $sorted_hvac \leftarrow$ sort again based on ON/OFF (ON first)

At t_1 : however, the regulation signal becomes two units, so both HVAC units need to be turned ON to track the signal. Algorithm 1 can turn on HVAC A, and keep also the HVAC B ON to meet the requirement. But, Algorithm 2 which just turned OFF HVAC B at t_0 cannot turn it back on just yet, so it cannot meet the regulation signal requirement. Both algorithms can track the signal from t_2 onwards. Thus, for Case II, Algorithm 1 acts optimally, whereas Algorithm 2 does not.

This counter example shows that the same decision made at time t_0 can become either: (i) the only choice to make the aggregated HVAC power follow the regulation signal most closely; or (ii) the choice that prevents the aggregated HVAC power from following the regulation signal most closely depending upon how the regulation signals change from t_1 onwards. This counter example can serve as a counter example for the case of hundreds of HVACs as well, by simply assuming that the half of those HVACs behave like HVAC A and the other half behave like HVAC B. It should be noted that, although this counter example might seem to be simplified, any counter-example in the problem space is sufficient to prove the non-existence of the universal optimal algorithm because, to be universally optimal, such algorithm should be optimal for all conceivable situations of the problem space.

Hence, a preemptive optimal regulation algorithm cannot exist. As such, only heuristic algorithms can be designed that can perform decently well in a range of situations. The proposed heuristic algorithms are discussed in the next section.

IV. THE PROPOSED SOLUTION

Two heuristic algorithms were developed to deal with aggregated HVAC control following regulation signals, namely: Greedy Algorithm (GA) and Lazy Algorithm (LA). Note that because regulation signals vary every two seconds, our control time step was chosen to be two seconds. Effectively, the aggregated HVAC power is expected to vary every two seconds

following regulation signals, while at the same time, allowing no HVAC compressor to cycle faster than their minimum ON/OFF times of two/three minutes, respectively.

A. Greedy Algorithm (GA)

Rooted on the juggling algorithm presented in our previous work [21], GA was developed to track dynamic two-second regulation signals by turning on the HVACs with the earliest *time-to-boundary*, incorporating the minimum HVAC ON/OFF time constraints and taking into account customer comfort. This algorithm is detailed below.

The sorting by time-to-boundary in line 6 ensures that the HVAC with the earliest need to turn ON is prioritized for turning ON during each iteration. Since the run-time complexity of sorting is log-linear [27], the algorithm has a log-linear time complexity with respect to the number of HVACs.

The *time-to-boundary* ($B_n^{t_k}$) of HVAC_{*n*} (HVAC of house *n*) at a given time t_k is defined as the time it takes for the temperature to hit the upper limit if the HVAC remains OFF. The $D_n^{t_k}$ in the above algorithm is the time by which the *time-to-boundary* of HVAC_{*n*} is delayed (increased) when HVAC_{*n*} is turned ON during a control period. It depends on house properties and HVAC capacity. Its numerical value can be determined by calculating the difference between two *time-to-boundary* values obtained using: one—chosen at the current temperature, and the other—made equal to the temperature attained when the HVAC cools the building for one control period [21]. The control period is the control time step (two seconds) if the HVAC is already ON, or is the minimum HVAC ON time (two minutes) if the current status of the HVAC is OFF.

B_n^{\max} is the maximum allowable *time-to-boundary* for HVAC_{*n*}, which is a proxy for the lower temperature limit.

The regulation MW level to be met ($D_L^{t_k}$) at time step t_k is calculated as:

$$D_L^{t_k} = \text{reg}^{t_k} * \text{RegMW}^{hr} / 2 + \text{midPower}^{hr} \quad (5)$$

where,

reg^{t_k} Current regulation signal value (varies from -1 to $+1$) at time t_k

hr Current hour, $\text{int}(t_k/3600)$

RegMW^{hr} Regulation capability being delivered for the current hour (hr)

midPower^{hr} Midpoint power level of the current hour

Determination of RegMW^{hr} and midPower^{hr} is discussed in Section IV-C.

The function $\text{mustRun}(\text{HVAC}_n)/\text{mustNotRun}(\text{HVAC}_n)$ determines if HVAC_{*n*} must run or must not run in the current time step. It returns true if any of the following conditions are met: (i) HVAC_{*n*} was recently turned ON/OFF less than its minimum ON/OFF time; (ii) if HVAC_{*n*} is turned ON/OFF right now but turning it OFF/ON for at least its minimum OFF/ON time would result in the indoor temperature hitting the upper/lower boundary; and (iii) if HVAC_{*n*} is turned OFF/ON, but letting it remain OFF/ON for an additional time step would make the indoor temperature hit the upper/lower boundaries.

This algorithm is named ‘greedy’ because it greedily prioritizes HVACs with the earliest *time-to-boundary* for being candidates to turn ON.

B. Lazy Algorithm (LA)

In this study, LA was developed as a modification to the GA. In particular, instead of prioritizing HVACs with their earliest *time-to-boundary*, it prioritizes maintaining the state of HVACs (ON or OFF) so as to minimize the state changes. It tends to avoid the job of changing HVAC states for as long as possible, so it is named lazy algorithm. The detail of this algorithm is outlined below:

The LA is mostly the same as the GA except that the *sorted_hvac* is stable-sorted again based on the current ON/OFF status of HVACs. This results in HVACs that are already ON to be prioritized for remaining ON (and consequently the HVACs which are OFF are prioritized for remaining OFF) in the current time step. If the current power requirement is unmet by the currently ON HVACs, only then OFF HVACs are considered for turning ON. Within the ON and OFF groups, the decisions remain prioritized by their *time-to-boundary*.

It can be noted that the problem of inadvertent synchronization of HVAC typical of setpoint change based control does not happen with either of the algorithms because, the state of the HVAC is explicitly controlled by the algorithm as per the heuristic as opposed to just controlling the thermostat setpoint.

C. Determining midPower^{hr} and RegMW

midPower^{hr} and RegMW are determined as follows:

1) *Determining the Midpoint Level (midPower^{hr})*: One of the major constraints is to maintain the indoor temperatures of all participating houses within the upper and lower temperature boundaries. The indoor temperature is the function of the cooling energy expended by an HVAC. If the total hourly energy expenditure during a regulation period remains the same as the energy expenditure without regulation, the average indoor temperature of the house is expected to be about the same level with or without regulation. Studying PJM’s historical RegD signals shows that although the mean value varies hour to hour on average in a year, the mean is centered around zero. This implies, if the midpoint level is chosen to be same as the hourly average HVAC power consumption (without regulation), the same amount of energy is consumed with regulation as without regulation. Hence, the average indoor temperature can be more or less maintained to be around the same value without letting them drift.

Thus, in this study the midpoint power level was chosen as: $\text{midPower}^{hr} = \text{basePower}^{hr}$. Where, basePower^{hr} is the average power consumption of all HVACs for hour hr without regulation.

2) *Determining the RegMW*: RegMW should be selected such that the total regulation credit (TRC), which is proportional to $\text{PS} * \text{RegMW}$ as per (3), is maximized. Since the midpoint power is fixed at basePower^{hr} , the maximum value of RegMW, $\text{RegMW}_{\max} = \min(2 * \text{midPower}^{hr}, 2 * (\text{maxPower}_{hr} - \text{midPower}^{hr}))$, where maxPower_{hr} is the total power of all the HVACs combined.

As RegMW is decreased from its maximum value, PS can increase, thereby having the potential for the credit to increase. A binary search algorithm, similar to that used in [21] to

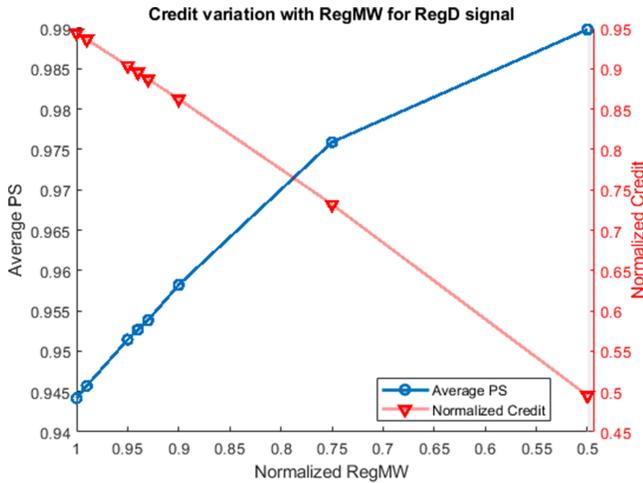


Fig. 3. Average PS and regulation credit variations with RegMW for RegD.

determine best demand limit, could be used to determine the optimal RegMW for each hour. This was accomplished by iterating the RegMW value and using a simulation model to determine the performance score and calculating the TRC at each iteration step. However, the simulation study –carried out by varying the RegMW down from the maximum towards its 50% value for a various assortment of real world regulation signals and for various hours of the day– shows that the PS does not increase fast enough or at all to compensate for the reduced RegMW. Hence, the maximum credit is always available at the maximum RegMW. This is illustrated in Fig. 3 for various values of RegMW normalized with respect to $RegMW_{max}$ for the simulation run of 100 houses using LA, as GA and LA exhibit similar properties.

D. Strength and Weakness of the GA and LA

As GA tends to turn ON the HVACs with the earliest *time-to-boundary*, GA always favors to keep *time-to-boundaries* of different houses close and as much away from zero as possible. Hence, it results in frequent HVAC state changes. If there is an abrupt change in a regulation signal at any given time, many HVACs are likely to be locked out (because of recent state transitions) and cannot immediately respond to the signal. However, because the *time-to-boundaries* are concentrated and away from zero, HVAC states would likely be maintained for considerably long time compared to LA before the HVACs are forced to turn ON/OFF to prevent *time-to-boundary* from becoming zero (i.e., the temperature hitting the boundary).

On the other hand, LA lets HVACs maintain their states unless the temperature comfort constraints are to be violated. This minimizes unnecessary state transitions, and reserves the opportunity to abruptly change states if required in a short notice as per the regulation signal. However, it lets the *time-to-boundaries* to disperse and float near zero or *maximum-time-to-boundaries*. Hence, should an abrupt change of large magnitude occur in the regulation signal (this will require ON/OFF responses from most of the participating HVACs), because the *time-to-boundaries* are already near the limits, this

algorithm will not be capable of delivering that level of power for an extended period of time.

Since each algorithm has its weakness and strength, the overall performance will depend upon the nature of the regulation signal. Next section describes the simulation study and the performance of each algorithm tested using real-world regulation signals. In addition, the impacts of temperature comfort range, regulation signal type and signal delay on the PS and TRC are discussed.

V. SIMULATION STUDIES AND DISCUSSION

Simulation studies for a collection of 100 houses were conducted using SimPy [28] – a python based discrete event simulation library, using the ETP model [21], [22] as the thermal models of residential houses, similar to the house_e model in GridLAB-D [24], [25]. House floor areas and other thermal parameters were randomized around typical values (e.g., floor area avg of 2200 sq ft). Typical meteorological year outdoor temperature and solar insolation data for Sterling, VA, were used. The simulation duration was set to 20 days starting on August 2nd, 2017 and the regulation signals for those 20 days were downloaded from PJM [31].

It is to be noted that the regulation signal has no correlation whatsoever between hours. The signal between any two hours in a certain day is as unrelated as the signal between the hours in a different year. So, as far as ensuring all varieties of regulation signal is concerned, we believe 20 days x 24 hours is sufficient. In a similar vein, the simulation was conducted for 20 days to introduce sufficient variation in the outdoor weather pattern. And to introduce variation in the house characteristics, the simulation was conducted on a collection of 100 houses with their thermal parameters randomly varied. So, in essence, each hour of those 480 hours of simulation serves a separate study.

Simulations without regulation services have been conducted with comfort range of $\pm 1^\circ\text{F}$ (a typical thermostat deadband) to determine the basePower for various base setpoints.

Because the minimum RegMW quantity that would qualify for participating in the wholesale market is 100kW, participation in the regulation services only occurs when $RegMW^{hr}$ is at least 100kW. In the following case studies, LA with the comfort range of $\pm 2^\circ\text{F}$, the base setpoint of 77°F (based on [32]), RegD signals with no communication delay were used, unless otherwise noted. While PS is shown for each hour in Section V-A, averaged PS over the 20-day regulation services is shown in other Subsections.

A. GA vs LA vs Random Scheduling

The variation of PS for all hours during the simulation when regulation was provided is plotted in the Fig. 4 for both GA and LA. As shown, for almost all of the cases, LA outperforms GA by a good margin. The average PS for GA was 89.37% while it was 94.42% for LA. For the sake of comparison, a simulation run was also conducted with the random scheduling algorithm (RA), where the algorithm randomly picks which HVACs to turn ON/OFF from the pool of eligible

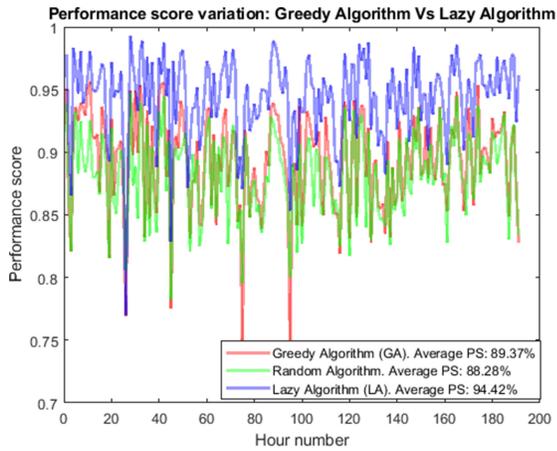


Fig. 4. PS for various regulation cases with different algorithms.

HVACs, similar to that proposed in [17], [18]. The average PS for the RA was 88.28%, so both GA and LA proposed in this paper perform better than RA.

Performance of GA and LA during a typical day is depicted in Fig. 5. Fig. 5(a) shows the performance of LA that makes the aggregated household load follow the regulation signal during hours 9 to 10 with PS of 94.78%. Under the same condition, GA in Fig. 5(b) performs slightly poorly with PS of only 90.36%. Although there were a few hours during which GA performed slightly better than LA, LA was found to perform much better than GA for almost all the cases. Hence, LA was used for the remaining analysis in determining the impacts of setpoints, thermal comfort constraint, communication delay and regulation signal type on regulation services. It can be noted in Fig. 5(a) and Fig. 5(b) that the indoor temperatures of the houses are kept within $\pm 2^\circ\text{F}$ of the 77°F setpoint in both cases.

B. Effects of Using Different Base Set Points

The base setpoint, i.e., the midpoint value of the temperature comfort range of the houses, was chosen to be 77°F in this study. A simulation sweep was conducted with the base setpoint of 75°F , 76°F , 77°F , 78°F , and 79°F to study its effect on PS, $\text{RegMW}_{\max}^{\text{hr}}$ and consequently the regulation credit (TRC). The plot in Fig. 6(a) shows that as the base set point is lowered from 77°F , the average $\text{RegMW}_{\max}^{\text{hr}}$ increases. This is understandable because it requires more energy to maintain a lower average temperature, and hence the $\text{basePower}^{\text{hr}}$ needs to increase, which allows for higher $\text{RegMW}_{\max}^{\text{hr}}$. The average PS remains more or less the same, so the regulation credit increases with lower set points, though not by much.

C. Effects of Varying the Comfort Range

A sweep of comfort range from 1°F to 5°F was conducted to study its impact on PS (and consequently directly on TRC). The result in Fig. 6(b) shows that there is a sharp increase in PS when going from comfort range of 1°F to 2°F , but it quickly saturates. Hence, $\pm 2^\circ\text{F}$ seems to be a reasonable tradeoff between comfort and performance.

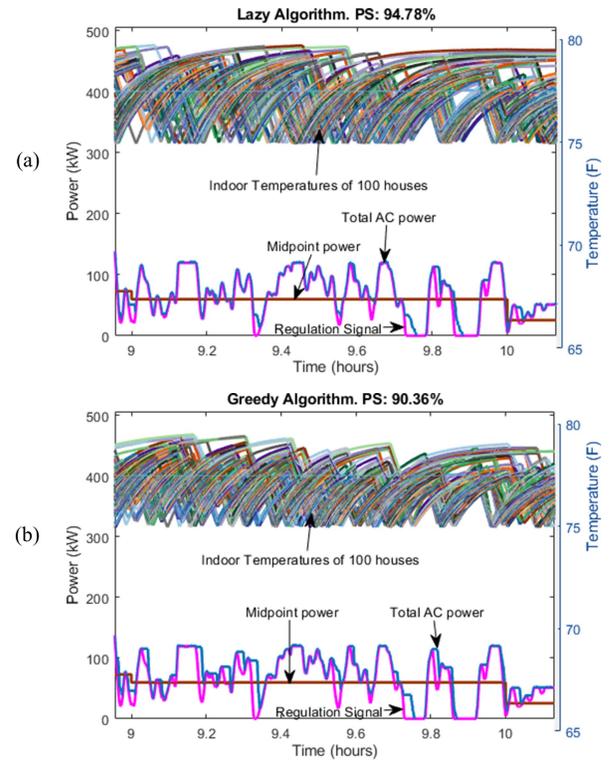


Fig. 5. A typical example day difference in PS: (a) LA performing better than GA; and (b) GA performing poorer than LA.

D. Effects of Communication Time Delay

So far, we have assumed that the DR aggregator can change the ON/OFF status of AC compressors with a negligible time delay. However, time delays always exist when communicating with smart thermostats, i.e., sending control commands to adjust HVAC set points in real-world implementation. The impact on PS when introducing various communication delays is summarized in Fig. 6(c), showing the resulting PS with varying delay from 0 to 20 seconds. As expected, PS keeps degrading as the delay increases, but it still remains quite high even up to five (5) seconds of communication delay. As such, LA appears to perform quite well even in real-world scenarios with delays.

E. Performance Analysis With RegA Signals

So far, RegD has been used because, unlike RegA, it is expected to be zero centered. As such, for RegA it is expected that temperatures of participating houses to be saturated at near their upper/lower comfort boundaries more often, and the PS to be not as good as with RegD. Fig. 6(d) shows PS while using RegA signals with various values of RegMW, normalized with respect to RegMW_{\max} . PS does improve, in general, when RegMW is reduced considerably. However, the PS does not improve fast enough to compensate for the reduced RegMW to improve the regulation credit. As such, just as in the case of RegD, the maximum credit is still obtained at RegMW_{\max} . But this credit is about 0.7434 units compared to the 0.944 units in the case of RegD (Fig. 3), which is 21.2% less. As such, aggregated HVAC control is much more suitable for RegD signal

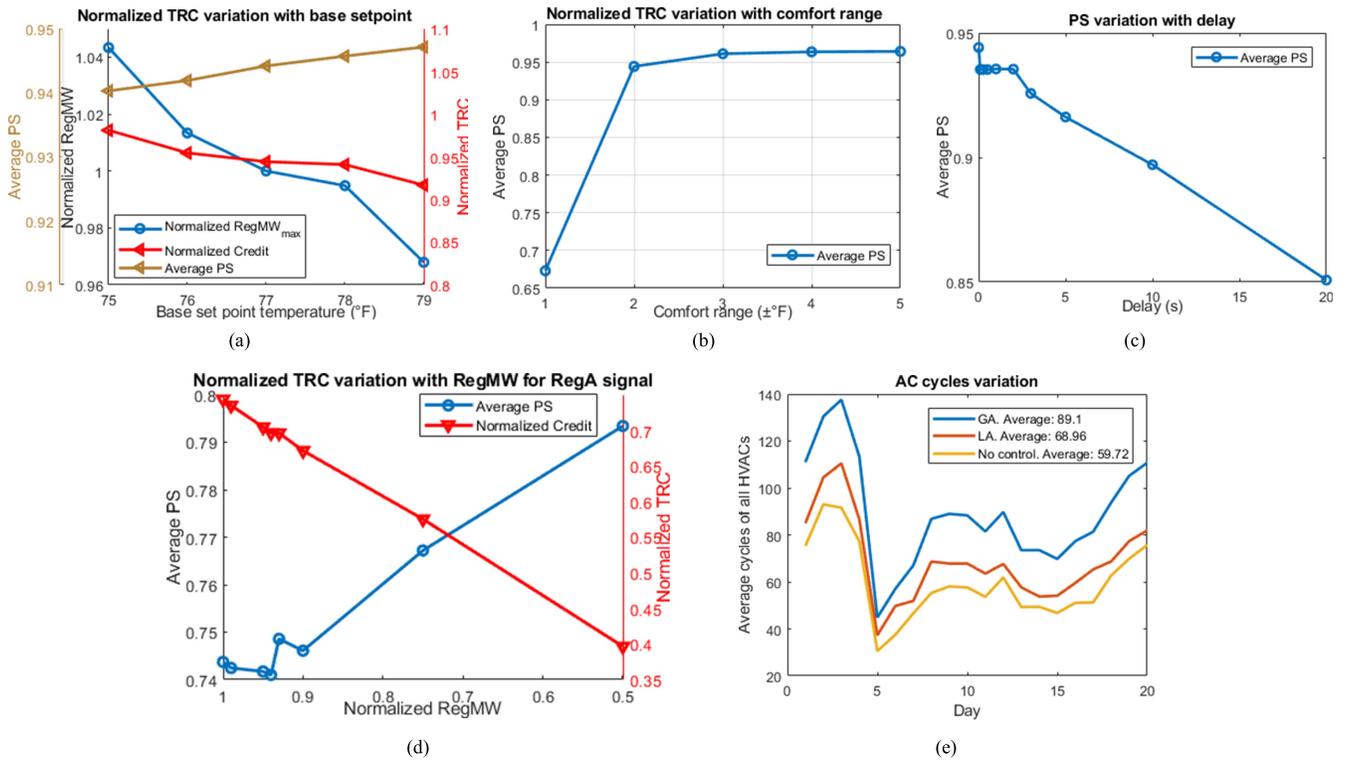


Fig. 6. (a) Impact of base set point on regulation credit; (b) Impact of comfort range on PS; (c) Impact of time delay on PS; (d) Average PS and regulation credit variations with RegMW for RegA signals; and (e) Impact of regulation on AC cycling.

than RegA. The peculiar decrease and increase of PS around the normalized RegMW of 0.95 and higher (see Fig. 6(d)) is because of the peculiarity of the regulation signal. Serving a higher RegMW in general is harder because it requires turning ON/OFF a large number HVACs which might not always be possible because of the temperature constraints. However, sometimes, the opportunity to turn ON greater number of HVACs can become useful to keep the temperatures within constraints, if in the preceding time period, they have been turned OFF for a long period of time.

F. Impact on the Number of AC Cycling

One concern while performing DR using HVAC is if the control results in an increasing number of AC cycling, thereby reducing the service life of the compressor. Fig. 6(e) shows the average number of daily cycles of all ACs for each of the 20 days, which varies based on the control algorithm used. It can be seen that GA always results in a substantial increase in the average number of AC cycling, but LA results in a much less increase in the AC cycling. This is understandable since LA tends to avoid cycling as much as possible (hence lazy).

G. Summary

Results of the simulation studies have been summarized in the Table I. The light blue tags highlight the parameters that are swept, and the orange tags highlight the values influenced by that sweep. The dark orange highlights the best result of each sweep.

TABLE I
RESULT SUMMARY

Ind ex	Alg orith m	RegM W	Base Set point	Com- fort Range	Delat y	Reg Sig- nal	PS	TR C
1	LA	1.0	77	2	0	D	.944	.944
2	GA	1.0	77	2	0	D	.894	.894
3	LA	1.0	77	2	0	D	.944	.944
4	LA	0.9	77	2	0	D	.958	.862
5	LA	0.75	77	2	0	D	.976	.732
6	LA	1.0	77	2	0	A	.743	.743
7	LA	0.9	77	2	0	A	.746	.672
8	LA	0.75	77	2	0	A	.767	.575
9	LA	1.043	75	2	0	D	.940	.981
10	LA	1	77	2	0	D	.944	.945
11	LA	0.968	79	2	0	D	.947	.916
12	LA	1.0	77	1	0	D	.673	.673
13	LA	1.0	77	2	0	D	.944	.944
14	LA	1.0	77	3	0	D	.961	.961
15	LA	1.0	77	5	0	D	.964	.964
16	LA	1.0	77	2	0	D	.944	.944
17	LA	1.0	77	2	5	D	.916	.916
18	LA	1.0	77	2	10	D	.897	.897

Rows 1-2 indicates that LA has superior performance to GA with RegD. Rows 3-8 indicate that the performance of LA is better with RegD than RegA, and that PS and TRC decrease with the decrease in RegMW. Rows 9-11 illustrate the impact of base set point variation. Rows 12-15 show the impact of comfort range variations. And, rows 16-18 are for communication delay variation. The case with the 75°F base set point with comfort range of $\pm 5^\circ\text{F}$ and no communication delay gives the maximum TRC for each of those sweeps.

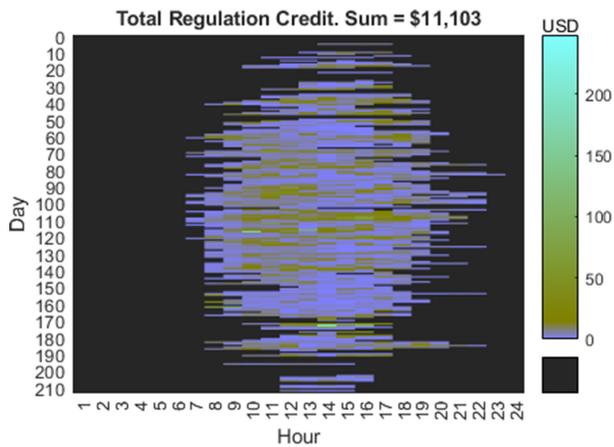


Fig. 7. Total regulation credit in dollars. Black areas are hours when the total regulation capacity was less than 100kW and no regulation is provided.

H. Regulation Credit Estimate for 100 Houses

The final simulation run was conducted using LA, assuming a realistic two-second communication delay for smart thermostat control. This was to determine the dollar amount an aggregator can expect to earn per year by selling regulation services from an aggregation of 100 houses. The study period was from April 1st to October 31st (214 days, when the AC was in operation). Regulation services were assumed to be provided during all hours when at least 100kW of RegMW could be provided. The CCP, PCP and MR for those time period were downloaded from PJM. Using (3), TRC was calculated for all those hours. TRC is shown in Fig. 7 as the heatmap plot.

The figure indicates that the credits are low for most hours, but occasionally, credits can be higher than \$100 for some hours. These occur when the CCP and PCP become exceptionally high due to market dynamics. In this study, the total credit adds up to \$11,103, effectively averaging out to be about \$111 per house per annum.

VI. CONCLUSION

The LA proposed in this paper was found to be able to provide up to 94.2% performance score while following real-world RegD signals from PJM, compared to 89.37% provided by GA. Also, HVAC control was found to be more suitable for RegD signals, providing about 27% more credit than RegA signals. This was understandable since RegD was zero centered and did not require as much energy storage capability in the resource as RegA. An example simulation with 100 houses for the whole year conducted using the market clearing price data from PJM showed that a total regulation credit of up to \$11,103 per annum could be obtained. The proposed algorithm hence potentially serves as a practical tool for DR aggregators to explore the market for regulation services through an aggregated control of residential HVACs. In addition, we showed that the optimal algorithm for regulation was not possible and the heuristic algorithm presented in this paper was shown to perform well in a range of real-world situations. But it might be possible to find approximate

optimal solution, and this could be a topic of further research.

REFERENCES

- [1] *PJM 101 the Basics*. Accessed: Feb. 19, 2018. [Online]. Available: <http://www.pjm.com/training/~media/2F50B2D760BF44F1BC7C5A2A42EA2D2A.ashx>
- [2] N. A. Sinityn, S. Kundu, and S. Backhaus, "Safe protocols for generating power pulses with heterogeneous populations of thermostatically controlled loads," *Energy Convers. Manag.*, vol. 67, pp. 297–308, Mar. 2013.
- [3] N. Lu and Y. Zhang, "Design considerations of a centralized load controller using thermostatically controlled appliances for continuous regulation reserves," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 914–921, Jun. 2013.
- [4] M. Vanouni and N. Lu, "Improving the centralized control of thermostatically controlled appliances by obtaining the right information," *IEEE Trans. Smart Grid*, vol. 6, no. 2, pp. 946–948, Mar. 2015.
- [5] M. Ali, M. Humayun, M. Degefa, A. Alahäivälä, M. Lehtonen, and A. Safdarian, "A framework for activating residential HVAC demand response for wind generation balancing," in *Proc. IEEE Innov. Smart Grid Technol. Asia (ISGT ASIA)*, Bangkok, Thailand, 2015, pp. 1–6.
- [6] M. T. Muhssin, L. M. Cipcigan, N. Jenkins, S. Slater, M. Cheng, and Z. A. Obaid, "Dynamic frequency response from controlled domestic heat pumps," *IEEE Trans. Power Syst.*, vol. 33, no. 5, pp. 4948–4957, Sep. 2018.
- [7] J. Hu *et al.*, "Load following of multiple heterogeneous TCL aggregators by centralized control," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 3157–3167, Jul. 2017.
- [8] K. Ma, C. Yuan, J. Yang, J. Li, S. Hu, and Y. Wang, "Controller design and parameter optimization of aggregated thermostatically controlled loads for frequency regulation," in *Proc. 35th Chin. Control Conf. (CCC)*, Chengdu, China, 2016, pp. 10035–10040.
- [9] C. Li, Y. Chen, F. Luo, Z. Xu, and Y. Zheng, "Real-time decision making model for thermostatically controlled load aggregators by natural aggregation algorithm," in *Proc. IEEE Int. Conf. Energy Internet (ICEI)*, Beijing, China, 2017, pp. 279–284.
- [10] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop, "Experimental evaluation of frequency regulation from commercial building HVAC systems," *IEEE Trans. Smart Grid*, vol. 6, no. 2, pp. 776–783, Mar. 2015.
- [11] I. Beil, I. Hiskens, and S. Backhaus, "Frequency regulation from commercial building HVAC demand response," *Proc. IEEE*, vol. 104, no. 4, pp. 745–757, Apr. 2016.
- [12] B. M. Sanandaji, T. L. Vincent, and K. Poolla, "Ramping rate flexibility of residential HVAC loads," *IEEE Trans. Sustain. Energy*, vol. 7, no. 2, pp. 865–874, Apr. 2016.
- [13] M. Ghanavati and A. Chakravarthy, "Demand-side energy management by use of a design-then-approximate controller for aggregated thermostatic loads," *IEEE Trans. Control Syst. Technol.*, vol. 26, no. 4, pp. 1439–1448, Jul. 2018.
- [14] H. Hao, T. Middelkoop, P. Barooah, and S. P. Meyn, "How demand response from commercial buildings will provide the regulation needs of the grid," in *Proc. 50th Annu. Allerton Conf. Commun. Control Comput. (Allerton)*, Monticello, IL, USA, 2012, pp. 1908–1913.
- [15] K. Dehghanpour and S. Afsharnia, "Electrical demand side contribution to frequency control in power systems: A review on technical aspects," *Renew. Sustain. Energy Rev.*, vol. 41, pp. 1267–1276, Jan. 2015.
- [16] M. Liu and Y. Shi, "Optimal control of aggregated heterogeneous thermostatically controlled loads for regulation services," in *Proc. 54th IEEE Conf. Decis. Control (CDC)*, Osaka, Japan, 2015, pp. 5871–5876.
- [17] S. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proc. Power Syst. Comput. Conf. (PSCC)*, Stockholm, Sweden, 2011, pp. 1–7.
- [18] L. A. D. Espinosa, M. Almassalkhi, P. Hines, and J. Frolik, "System properties of packetized energy management for aggregated diverse resources," in *Proc. Power Syst. Comput. Conf. (PSCC)*, Dublin, Ireland, 2018, pp. 1–7.
- [19] Y.-Q. Bao, J.-Q. Ju, Q. Wang, and G. Ma, "On the accuracy of the state-queueing model for the thermostatically controlled loads," in *Proc. IEEE PES Asia-Pac. Power Energy Eng. Conf. (APPEEC)*, Xi'an, China, 2016, pp. 1990–1994.

- [20] M. Olama, T. Kuruganti, J. Nutaro, and J. Dong, "Coordination and control of building HVAC systems to provide frequency regulation to the electric grid," *Energies*, vol. 11, no. 7, p. 1852, Jul. 2018.
- [21] R. Adhikari, M. Pipattanasomporn, and S. Rahman, "An algorithm for optimal management of aggregated HVAC power demand using smart thermostats," *Appl. Energy*, vol. 217, pp. 166–177, May 2018.
- [22] PJM. (2017). *PJM Manual 11: Energy & Ancillary Services Market Operations*. Accessed: Feb. 22, 2018. [Online]. Available: <http://www.pjm.com/~media/documents/manuals/m11.ashx>
- [23] (2017). *PJM Manual 12: Balancing Operations*. Accessed: Feb. 22, 2018. [Online]. Available: <http://www.pjm.com/~media/documents/manuals/m12.ashx>
- [24] S. Benner. (2015). *Performance, Mileage and the Mileage Ratio*. Accessed: May 3, 2018. [Online]. Available: <https://www.pjm.com/~media/committees-groups/task-forces/rmistf/20151111/20151111-item-05-performance-based-regulation-concepts.ashx>
- [25] R. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Ph.D. dissertation, Dept. Aerosp. Mech. Sci., Princeton Univ., Princeton, NJ, USA, 1978.
- [26] Z. Taylor, K. Gowri, and S. Katipamula, *GridLAB-D Technical Support Document: Residential End-Use Module Version 1.0*, U.S. Dept. Energy, Washington, DC, USA, 2008.
- [27] *TimeComplexity—Python Wiki*. Accessed: May 3, 2019. [Online]. Available: <https://wiki.python.org/moin/TimeComplexity>
- [28] *SymPy: Discrete Event Simulation for Python*. Accessed: May 3, 2019. [Online]. Available: <https://simpy.readthedocs.io/en/latest/index.html>
- [29] *GridLAB-D House_e Model Source Code*. Accessed: Jun. 27, 2017. [Online]. Available: http://www.gridlabd.org/documents/doxygen/latest_dev/house_e_8cpp-source.html
- [30] *Residential Module User's Guide—GridLAB-D Wiki*. Accessed: Jun. 26, 2017. [Online]. Available: http://gridlab-d.shoutwiki.com/wiki/Residential_module_user%27s_guide
- [31] *PJM—Ancillary Services*. Accessed: Jan. 9, 2019. [Online]. Available: <https://www.pjm.com/markets-and-operations/ancillary-services.aspx>
- [32] *Thermal Environmental Conditions for Human Occupancy*, vol. 55, ANSI/ASHRAE Standard 55–2010, p. 42, 2010.



Rajendra Adhikari (M'19) received the Ph.D. degree in electrical and computer engineering from Virginia Polytechnic and State University (Virginia Tech) in 2018, focused on building energy management system and algorithms, and simulation framework for demand response. He is a Senior Researcher with National Renewable Energy Laboratory, researching in the areas of building energy efficiency, smart home control and the use of smart-contract-based block-chain technology for peer-to-peer energy trading, and transactive controls.



Manisa Pipattanasomporn (SM'16) is an Associate Professor with the Smart Grid Research Unit (SGRU), Chulalongkorn University, Thailand, and an Adjunct Faculty with the Advanced Research Institute, Virginia Tech, USA, where she has been with the Department of Electrical and Computer Engineering since 2006. In 2019, she joined SGRU, Chulalongkorn University, as a Research Associate Professor with the aim to develop an Internet of Things (IoT)-based smart campus. Her research interests include smart grid, smart home, smart building, demand response, renewable energy integration, electric vehicles, IoT, machine learning with particular applications on energy savings, building-level load forecasting, and blockchain for peer-to-peer trading of solar electricity.



Saifur Rahman (LF'13) is the Founding Director with the Advanced Research Institute, Virginia Tech, USA, where he is the J. R. Loring Professor of electrical and computer engineering. He also directs the Center for Energy and the Global Environment. He has published over 140 journal papers and has over 400 conference and invited presentations. He is an IEEE Millennium Medal Winner. He was the Founding Editor-in-Chief of *IEEE Electrification Magazine* and the IEEE TRANSACTIONS ON SUSTAINABLE ENERGY. He served as the Chair of the U.S. National Science Foundation Advisory Committee for International Science and Engineering from 2010 to 2013. He is the President of the IEEE Power & Energy Society (PES) for 2018 and 2019. He is a Distinguished Lecturer for the PES and has lectured on renewable energy, energy efficiency, smart grid, energy Internet, blockchain, and IoT sensor integration in over 30 countries.