ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Quantifying Flexibility of Residential Thermostatically Controlled Loads for Demand Response: A Data-driven Approach

Emre Can Kara[1,3], Michaelangelo D. Tabone[2], Jason S. MacDonald[3], Duncan S. Callaway[2,3], Sila Kiliccote[3]

[1] Carnegie Mellon University, Pittsburgh, PA
[2] University of California, Berkeley, CA
[3] Lawrence Berkeley National Lab, Berkeley, CA

October 2014

Part of the work described in this paper was coordinated by the Consortium for Electric Reliability Technology Solutions and was funded by the Office of Electricity Delivery and Energy Reliability, Transmission Reliability Program of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

BuildSys'14 November 5–6 2014, Memphis, TN,USA

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Quantifying Flexibility of Residential Thermostatically Controlled Loads for Demand Response: A Data-driven Approach

Emre Can Kara
Civil & Environmental Engineering
Carnegie Mellon University
Pittsburgh, PA 15213
eckara@cmu.edu

Michaelangelo D. Tabone
Energy & Resources Group
University of California, Berkeley
Berkeley, CA 94720
m.tabone@berkeley.edu

Jason S. MacDonald
Environmental Energy Technologies Division
Lawrence Berkeley National Laboratory
Berkeley, CA 94720
jsmacdonald@lbl.gov

Duncan S. Callaway
Energy & Resources Group
University of California, Berkeley
Berkeley, CA 94720
dcal@berkeley.edu

Sila Kiliccote
Environmental Energy Technologies Division
Lawrence Berkeley National Laboratory
Berkeley, CA 94720
skiliccote@lbl.gov

Abstract

Power systems are undergoing a paradigm shift due to the influx of variable renewable generation to the supply side. The resulting increased uncertainty has system operators looking to new resources, enabled by smart grid technologies, on the demand side to maintain the balance between supply and demand. This study uses a unique data set to estimate and validate models of demand response from residential thermostatically controlled loads (TCLs)—specifically, HVAC units—and quantifies the extent to which a population of TCLs can provide demand response (DR). We use measured temperature setpoints, internal temperatures, compressor cycling ratio and metered energy data collected from over 4200 homes in Texas during the summer of 2012. Using autoregressive moving average (ARMA) models for individual households, we investigate the instantaneous power shed, the duration of the power shed, steady state energy savings and total energy savings. Specifically, we provide insight into the dependency of household DR availability to the temperature setpoint schedule, outdoor air temperature and time of the day.

Categories and Subject Descriptors

D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*; H.4 [**Information Systems Applications**]: Miscellaneous; H.2 [**Data**]: Data mining

Keywords

Demand response, Thermostatically controlled loads, Inverse building model, Load shedding

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

BuildSys'14 November 5–6 2014, Memphis, TN,USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3144-9/14/11 ...$15.00
http://dx.doi.org/10.1145/2674061.2674082

1 Introduction

Renewable portfolio standards adopted in 29 states in the US require a certain percentage of their electricity generation to come from renewable generation sources. This requirement is intended to reduce the green house gas emissions caused by traditional fossil fuel-based power generation. However, increasing the penetration of renewable energy resources into the power grid calls for an increase in the availability of additional reserves in the power system to keep the supply and demand in balance [7, 6]. Automated demand response (ADR), in which load adjustments are made by automated actions (in contrast to price-based demand response, which relies on consumer actions), is a promising resource that can provide the required flexibility and eliminate the need to hold fossil fuel-based generators in reserve [7, 6, 15].

In recent years, thermostatically controlled loads (TCLs) such as heating, ventilation, and air conditioning (HVAC) systems, refrigerators, and water heaters have been garnering interest in the research community as ADR resources. This is partly because of their wide availability and partly because of the thermal storage that they possess; this storage allows them to be turned on and off for an undetermined amount of time without affecting the performance of the thermostatic control.

Direct load control (DLC) of a population of residential TCLs has been studied by various researchers [10, 11, 12, 17, 2, 8, 19]. A linear time-invariant representation and a Markov decision process (MDP)-based representation of a heterogeneous TCL population using state-bin transition models is given in [14] and in [10], respectively. A priority stack-based algorithm to provide ancillary services is presented by [8]. The authors of [18] use household refrigeration units that are modified with additional thermal storage to provide peak shaving.

In addition, the quantification of TCLs' capability to provide DR has also been investigated by various researchers through bottom-up simulation methods, where simulation model parameters are based on educated guesses [8, 13]. Studies investigating the impacts of ADR on the power grid level include, but are not limited to, [4, 3, 9, 1].

However, in existing studies, quantification of the resources are

based on bottom-up simulations using plausible but not rigorously identified parameters. To the best of our knowledge, [5] is the only work that uses real electricity consumption data to estimate models that can then be used to compute DR potential. However, the dataset used in [5] consists of whole-building smart meter data, so a great deal of the modeling effort (and uncertainty) is due to the challenge of disaggregating HVAC demand from whole-buildings' demand.

This study's central contributions are (1) a method to use direct HVAC measurements to estimate DR potential, and (2) initial estimates of DR potential from a large aggregation of residential loads using this method. We use a unique dataset in which HVAC consumption, temperature setpoint and measured indoor temperature were recorded for over 4200 households. This allows us to sidestep the issue of disaggregation encountered in [5]. Using ARMA models trained based on historical data from individual households, we investigate the impacts of varying exogenous parameters on flexibility metrics such as instantaneous power shed, the load shed duration, steady-state energy savings and total energy savings. We investigate the dependency of household DR availability to the temperature setpoint schedule, outdoor air temperature and time of day.

We find that the constant exogenous input assumption made commonly in the literature (e.g. [10, 11]) yields to a significant over-estimation of the number of loads that can provide a requested shed duration under certain cases. For the cases with a larger setpoint adjustment, the steady-state energy savings are higher for both weekday and weekend profiles.

The paper is organized as follows: Section 2 describes the dataset used in this study. Section 3 introduces the ARMA model built on [16] and evaluates the performance of this model for a single household. Section 4 uses the model proposed in Section 3 to estimate the DR flexibility, and discusses the sensitivity of these estimations to exogenous inputs. Section 5 presents the results of the DR flexibility estimation, and finally, conclusions are drawn and future work is suggested in Section 6.

2 Dataset

The dataset used in this study is obtained from 4297 households located in Texas. The data include temperature setpoint, indoor temperature, compressor cycling ratio and energy consumption measurements sampled every 5 minutes during the summer of 2012. The compressor cycling ratio is defined as the ratio of the time that the HVAC compressor is ON within a sampling period to the sampling period. In addition, hourly weather data obtained from weather stations in the Texas area closest to the households are used to capture the external temperature. The energy consumption measurements are converted into average power measurements assuming constant power use within each sampling interval. A solar earth geometry model was used to find the solar insolation on a horizontal surface at each home's approximate location at each time of day, assuming no cloud cover. Specific data fields used in this study along with their sampling rate and resolution information are given in Table 1. Data was only collected from homes with communicating thermostats, and thus are not a representative sample of all homes in the service territory. However, if a demand response program is to only control homes which have communicating thermostats (a very likely case), the sample does represent homes that are currently capable of participating.

A daily snapshot of each data field in Table 1 for a sample household is given in Figure 1.

Table 1. Data field characteristics

Data Field	Resolution	Sampling Rate
Temperature setpoint, $T_{set}(t)$	1 $[^oF]$	5 mins
Internal Temperature, $T_{int}(t)$	1 $[^oF]$	
Duty Ratio, $d(t)$	0.001	
Average Power, P_{avg}	0.01 $[W]$	
External Temperature, $T_{ext}(t)$	0.1 $[^oC]$	1 hour
Solar Insolation, $\phi_{sol}(t)$	0.01 W/m^2	

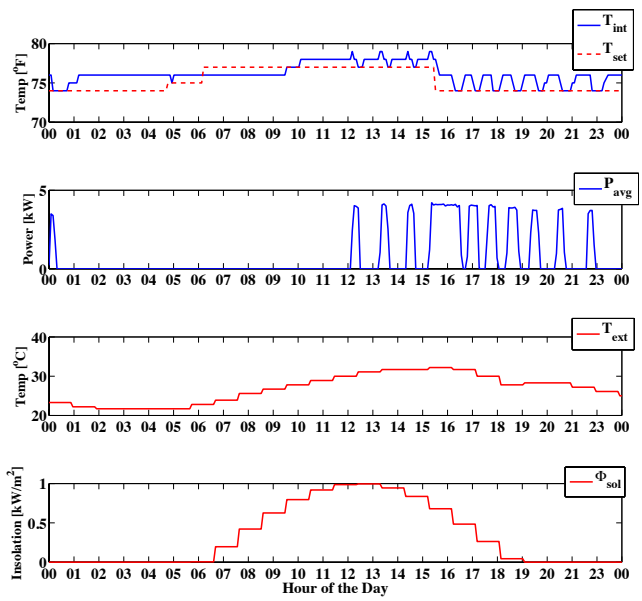

Figure 1. A daily snapshot of each data field used in this study for a sample household.

3 Estimating Thermal Models

We fit thermal models to each building monitored in our dataset with the objective of simulating transient dynamics of internal temperature and power consumption during an ADR event. Two important parameters of our model were apparent from the data without detailed statistical models: average temperature setpoint of the thermostat, and instantaneous HVAC power consumption. We used linear regression to fit an ARMA model to account for appliance efficiencies, thermal masses, thermal resistances, interior heat gains and solar gains.

3.1 ARMA Model

Eq. (1) shows an ARMA model for the thermal dynamics of a building, reproduced from [16] by Ari Rabl. ARMA models allow predictions of internal temperature to be dependent on both coincident and prior (a.k.a. "lagged") values of physical inputs such as cooling energy, outdoor temperature, and solar isolation. In [16], Rabl shows that ARMA models are the most general possible building model; he also shows that models with more direct physical analogies are reducible to ARMA models (e.g. differential equation models and networks of thermal resistances/capacitances). In the remainder of this section, we fully describe the ARMA model described in Eq. (1). We also detail how we constrained the parameters such that physical principles hold at steady state.

In Eq. (1), $T_{int}(t)$ is the room temperature of the house at time t, $T_{ext}(t)$ is the outdoor ambient temperature, $\phi_{aux}(t)$ is the average auxiliary power over the interval (from an air conditioner or a heater), and $\phi_{sol}(t)$ is solar insolation. The summation terms include time-lagged readings of each variable; N_x is the number of

lagged readings included for the variable defined by the subscript x.

Model coefficients are denoted as $a_x(j)$ where the subscript, x, indicates the corresponding variable and j indicates the timing of the lag (in number of readings prior to present). a_{occ} is an intercept term which is analogous to a constant internal heat gain resulting from occupants and devices.

$$\sum_{i=0}^{N_{int}} a_{int}(i)T_{int}(t-i) = a_{occ} + \sum_{i=0}^{N_{out}} a_{ext}(i)T_{ext}(t-i) + \sum_{i=0}^{N_{sol}} a_{sol}(i)\phi_{sol}(t-i) + \sum_{i=0}^{N_{aux}} a_{aux}(i)\phi_{aux}(t-i) + \varepsilon(t) \quad (1)$$

We applied two constraints to the ARMA model. The first constraint simply scales all coefficients such that $a_{int}(0) = 1$. The second constraint, shown in Eq. (5), ensures that the steady state thermal properties of the wall are consistent and was originally presented in [16]. We re-derive the second constraint for the reader in Eqs. (2)-(4). Eq. (2) shows the rearranged ARMA model from Eq. (1) where all variables are in steady-state; we define steady-state variables as $\bar{X} = X(t)\ \forall t$, and steady state coefficient as $\bar{a}_x = \sum_{i=0}^{N_x} a_x(i)$.

$$\bar{T}_{int}\bar{a}_{int} - \bar{T}_{ext}\bar{a}_{ext} = \bar{\phi}_{aux}\bar{a}_{aux} + \bar{\phi}_{sol}\bar{a}_{sol} + a_{occ} \quad (2)$$

Eq. (2) is arranged such that all values on the right-hand side represent heat flows into or out of the building, and values on the left represent indoor/outdoor temperatures. At steady state, heat transferred across the constant positive indoor/outdoor temperature differential ΔT should be equal and opposite to that transferred across its negative, $-\Delta T$, implying that the sum of the coefficient on internal temperatures are equal to the sum of coefficients on outdoor temperatures, shown in Eqs.(3)-(4). Eq. (5) shows the second constraint, which is a rearrangement of this equality condition.

$$L(\bar{T}_{int} - \bar{T}_{ext}) = \bar{\phi}_{aux}\bar{a}_{aux} + \bar{\phi}_{sol}\bar{a}_{sol} + a_{occ} \implies \quad (3)$$

$$L = \sum_{i=0}^{N_{int}} a_{int}(i) = \sum_{i=0}^{N_{out}} a_{ext}(i) \implies \quad (4)$$

$$a_{int}(N_{int}) = \sum_{i=0}^{N_{out}} a_{ext}(i) - \sum_{i=0}^{N_{int}-1} a_{int}(i) \quad (5)$$

Eq. (6) shows the linear model following the substitution of both constraints into Eq.(1) and solving for the latest internal temperature.

$$\begin{aligned} T_{int}(t) - T_{int}(t-N_{int}) = & \\ & \sum_{j=1}^{N_{int}-1} a_{in,j}(T_{int}(t-j) - T_{int}(t-N_{int})) \\ & + \sum_{j=0}^{N_{ext}} a_{out,j}(T_{ext}(t-j) - T_{int}(t-N_{int})) \\ & + \sum_{j=0}^{N_{sol}} a_{sol,j}\phi_{sol}(t-j) \\ & + \sum_{j=0}^{N_{aux}} a_{aux,j}\phi_{aux}(t-j) \\ & + a_{occ} + \varepsilon_t \end{aligned} \quad (6)$$

We fit the parameters shown in Eq. (6) for each household in the dataset using ordinary least square regression. We did not measure auxiliary heating/cooling energy directly; as a proxy, we substituted the average power consumption of the AC over the interval. This substitution implicitly assumes that the coefficient of performance (COP) of the AC is constant during the study period, thus the power consumption of the AC is directly proportional to the cooling energy provided by the AC. In actuality, the COP decreases as the difference between indoor and outdoor temperature increases. We explore the effects of this assumption in the conclusions.

For the purposes of this paper, the model order was chosen by visual inspection on a few houses: $N_{int} = 20$, $N_{ext} = 0$, $N_{aux} = 2$, and $N_{sol} = 0$. A more appropriate model selection process will be developed for future work. Twenty lagged internal temperatures were more than sufficient for most homes, as indicated by coefficients of long lags estimates to not be significantly different from zero. However, including superfluous lags did not affect the performance of the model and gave a conservative estimate of the model's complexity. Coefficients for lagged outdoor temperatures were difficult to identify and were left out of the model. Outdoor temperature was measured at roughly hourly intervals and then linearly interpolated; thus, there was not much variation at short lags. Including a few lagged values of auxiliary power was necessary, as there was often a noticeable delay between energy consumed by the HVAC and any response in room air temperature.

3.2 Thermal Properties of Buildings

Figures 2 and 3 show box plots of estimated properties of each building in the dataset. In these plots, boxes represent the interquartile range (25th to 75th percentile), and whiskers represent the 2.5^{th} and 97.5^{th} percentiles.

The **setpoints** of thermostats in our dataset varied frequently and regularly, presumably in an attempt to conserve energy by mitigating heat losses. Figure 2 presents distributions of the average setpoint of each thermostat at each hour ending compared to the average during the hour beginning at 3AM (HB 3); results are also stratified by weekday versus weekend. As shown, during the week, most households raised their set points by up to $4°F$ during the afternoon as compared to the night. Fewer than 5% of households lowered their set points by more than $5°F$, or raised it by more than $10°F$. During the weekend, households were less likely to raise their set points during the afternoon.

Panel A of Figure 3 shows the distribution of estimated rated power use for all HVACs in the dataset, which is defined as the average power used by HVACs during intervals when the compressor cycling ratio is 1. Most HVACs consume between 2 and 4 kW of electricity when running at full load, though estimates vary widely, from 0.5 to 6 kW.

Panel B of Figure 3 shows the distribution of **steady-state heat transfer coefficients**, h_{ss}, estimated at each home. h_{ss} represents the amount of additional power required by the air conditioner to increase the steady state indoor/outdoor temperature differential by $1\ °F$. Eq. (7) shows our formula for calculating h_{ss}, which is derived by rearranging Eq. (3):

$$h_{ss} = \frac{L}{\bar{a}_{aux}} = \frac{\sum_i a_{int}(i)}{\sum_i a_{aux}(i)} \quad (7)$$

Our definition of h_{ss} deviates from convention because we reference it to power consumed by the HVAC rather than cooling energy produced by the HVAC (which we do not measure). Thus, our coefficient accounts for (1) the thermal resistance of the building shell material, (2) the building shell area, and (3) the COP of the air conditioner. Most of these values are around $100\ W/°F$; however, they vary from less than 0 (obviously erroneous estimates) to $800\ W/°F$.

Panel C of Figure 3 shows the distribution of **effective occupant heat gain**, Q_{occ}, defined in Eq. (8) as the steady state, average

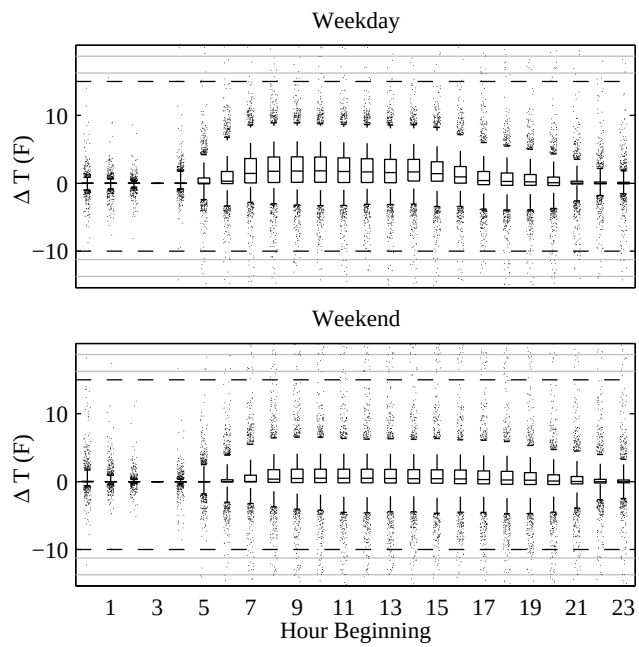

Figure 2. Average temperature setpoint for each house in the dataset, stratified by hour of day and by weekday-v-weekend. All setpoints are presented as deviations from the mean during HB3. Top panel shows average setpoints on weekdays; bottom panel shows weekends.

HVAC power required to exactly offset the occupant heat gain:

$$Q_{occ} = \frac{a_{occ}}{\bar{a}_{aux}} = \frac{a_{occ}}{\sum_i a_{aux}(i)} \tag{8}$$

Dividing Q_{occ} by h_{ss} gives the indoor/outdoor temperature differential sustained by only the occupant heat gain and no HVAC use.

Panel D of Figure 3 shows the distribution of **maximum solar heat gain**, Q_{sol}, which is defined in Eq (9) as the steady state average HVAC power required to offset solar heating gains at $1000W/m^2$:

$$Q_{sol} = \frac{1000\sum_i a_{sol}(i)}{\sum_i a_{aux}(i)} \tag{9}$$

As shown, distribution of these gains is centered at 0, meaning that the model parameters indicate that 50% of buildings derive *cooling energy* from solar insolation (or more specifically, from our solar-earth geometry model-derived estimates of clear-sky insolation). This result clearly indicates that these model parameters are not capturing the desired effect. This is likely a result of (1) shortcomings of using solar-earth geometry model data rather than solar observations (which would include cloud cover), (2) colinearity between the solar diurnal cycle and occupancy and (3) colinearity between the solar diurnal cycle and outside air temperature. We discuss these effects further in the conclusions.

Panel E of Figure 3 shows the distribution of **effective thermal capacity**, C_{eff}, which is defined in [16] as the amount of energy released by the building when reducing from a sustained internal temperature above steady state to the steady state. The value of C_{eff} is independent of all external variables as long as they are constant over time. We estimated this value by simulating the energy released by the building when returning from a steady temperature $1°F$ above steady state. The capacity of buildings in the dataset is typically between 1 and 2 $kWh/°F$. However, some buildings exhibit very large heat capacities.

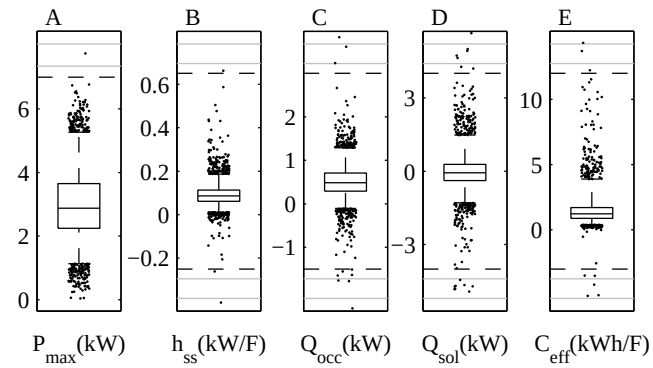

Figure 3. Distributions of estimated thermal properties of each house in the dataset.

3.3 Validation

We validate the model by predicting the response of internal temperature to large changes in set points that exist in the dataset. This validation metric is chosen to mimic the transient dynamics of an actual ADR event. Internal temperature was predicted based on knowledge of the internal temperature prior to the change (but not following), and the HVAC power consumption, outdoor temperature, and solar insolation throughout time period.

In order to be considered for validation, set point changes must be $4°F$ or larger and the set point must be constant for the two hours before and four hours after the change. We refer to a set point increase as a "load shed" and a set point decrease as a "load recovery," analogous to the beginning and end of an ADR event. There were 54,069 load sheds and 34,534 load recoveries with complete data in the dataset, and all were used for validation.

Figure 4 shows box plots of the model prediction errors at 1, 2, 3, 4, 5, and 6 hours following a set point change. The longer the time duration since the set point change, the greater the magnitude of the prediction errors. This result is expected because the set point change also marks the latest observations of actual internal temperature used to inform the prediction; predictions made six hours ahead of the last piece of information are more uncertain than those made one hour ahead.

Median prediction errors suggest that our model may be overestimating the effect of each building's thermal capacity. For load shedding events, the model tends to under-predict internal temperature directly following the change, signifying that the building is increasing in temperature more quickly than the model is predicting, and thus has less thermal capacity than predicted. However, for load recovery events, the model appears to be unbiased for the entire duration following the set point change.

The over-prediction of thermal capacity could be a result of an under-prediction of the effects of air conditioning and thermal gains on the internal temperature. In the conclusions, we suggest that this could be an effect of the low resolution of the internal temperature observations ($1°F$), as the rounding errors will be correlated with the effects of the cycling air conditioner.

4 Demand Response Potential

We use the ARMA model introduced in Eq. (6) to investigate the DR potential of an HVAC population under varying exogenous conditions. We simulate DR events by introducing a setpoint change in the thermostats, and we quantify the duration of the load shed $d_{i,t}^{shed}$

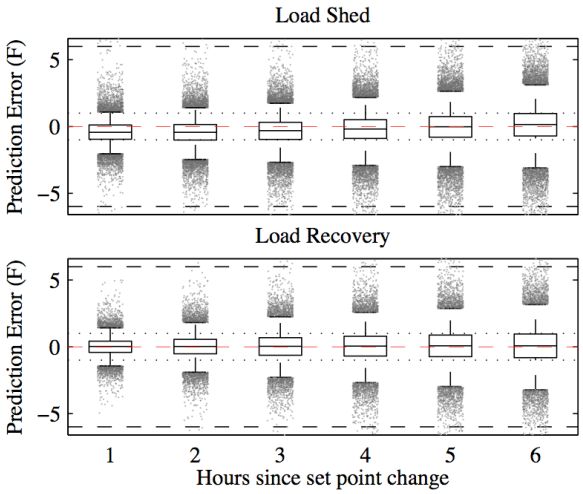

Figure 4. Model prediction errors following set point changes of greater than $4^{\circ}F$. Load shed events (top) are positive set point changes, and load recovery events (bottom) are negative set point changes.

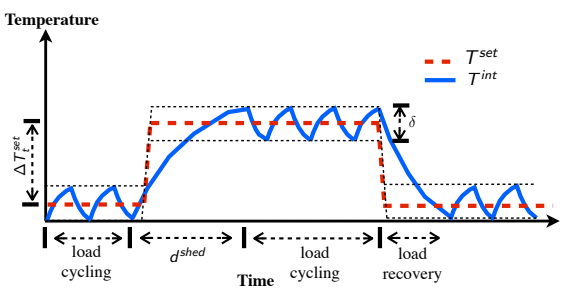

Figure 5. Event/control mechanism

given a setpoint change ΔT_t^{set} under different conditions. Furthermore, we investigate the behavior of the aggregate power consumption of the HVAC population and estimate potential energy savings due to shifting the HVAC loads to much cooler periods of the day.

Formally, we define an ADR event as follows: The internal set point T_{set} of the building is increased by a predetermined amount ΔT_t at time t. The duration of the load shed $d_{i,t}^{shed}$ for HVAC i at time t is defined as the time that it takes for $T_{i,t}^{int}$ to reach $T_{set} + \Delta T_t + \delta/2$, where δ is the thermostatic dead-band width. We simulate the behavior of each load as a 2-state load, whose thermal dynamics are governed by Eq. (6) and a thermostat. Specifically, we assume that the load turns OFF when the internal temperature hits the upper thermostatic bound (i.e. $T_{int} = T_{set} + \delta/2$) and vice versa. Figure 5 represents a typical load shed ADR event considered in this study and depicts the parameters defined above.

As expected, time varying setpoints affect a building's response to an ADR event. If the temperature setpoint is maintaining a high differential between indoor and outdoor temperature, an ADR event is expected to create a large instantaneous magnitude of power reduction, with a short duration (due to the high heat transfer over the differential). Conversely, if the temperature differential is small, the DR program should expect more modest instantaneous power reductions for a longer duration.

To capture this variation, we created various case studies using different setpoint profiles and exogenous input characteristics. We first investigate a simplistic scenario, where we assume all of the HVACs have the same temperature setpoint, the external air temperature is constant, and the insolation is zero. We then select a day to use as a representative day for the exogenous inputs such as: external temperature and insolation. We use average hourly setpoint schedules for HVAC loads obtained from the weekday and weekend measurements to obtain distributions of the $d_{i,t}^{shed}$ under varying conditions. The next section introduces the case studies in detail and discusses the assumptions made.

4.1 Case Studies

For all the case studies in this paper, we simulate a population of HVACs using the ARMA model given in Eq. (6). The simulation period and the thermostatic dead-band width is pre-determined and for all case studies t_{start} is 9AM, t_{end} is 11PM and δ is $1^{o}F$. At the beginning of each simulation, we assume the starting temperature for each HVAC, $T_{i,t_{start}}^{int}$, was randomly located within the thermostatic dead-band width. The HVAC statuses are also determined by observing the distribution of HVAC statuses at the starting hour from measured data based on the setpoint profile. During the simulation period, we assume that a single ADR event happens at t_{event}, with a predetermined setpoint change $\Delta T_{t_{event}}^{set}$. Since the Electric Reliability Council of Texas' (ERCOT) DR programs have varying dispatch durations ranging from 1 to 4 hours, we simulate the behavior of the aggregate load population with varying ADR event durations D_{DR}. Specifically, after a setpoint change of $\Delta T_{t_{event}}^{set}$ due to a DR even at time t_{event}, the setpoint of each HVAC is overwritten by $T_{i,t}^{set} + \Delta T_{t_{event}}^{set}$ for the duration of the ADR event, D_{DR}. When the event is over, all HVACs go back to their scheduled setpoint profile. Three sets of different case studies (CS-1, CS-2 and CS-3) are created to capture different exogenous input conditions. For each of these sets, we define 4 individual scenarios (e.g. A-D for CS-1) with varying D_{DR}, t_{event} and $\Delta T_{t_{event}}^{set}$ parameters. As seen in Table 2, the first set of case studies CS-1 focuses on estimating the DR potential of an HVAC aggregation under constant exogenous inputs. For the second set, we create an average weekday setpoint profile for all the HVACs, and assume that each HVAC is scheduled based on that. Similarly for the third set, we create an average weekend setpoint profile to differentiate the varying setpoint profiles due to change in occupancy patterns between weekdays and weekends.

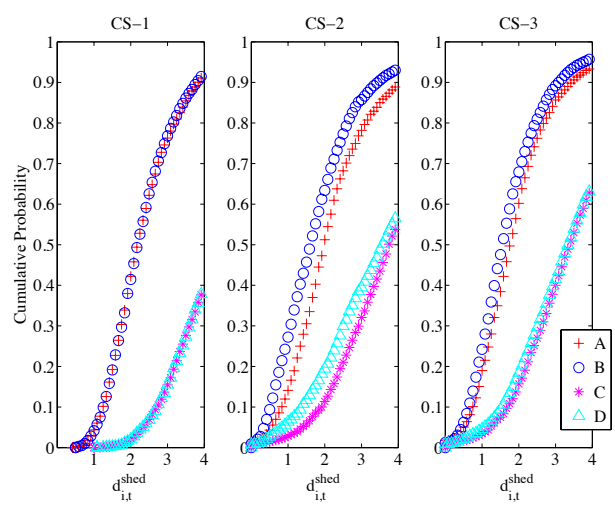

Figure 6. The cumulative probability distribution of d_{shed} obtained by simulating 4405 HVAC instances, available in the dataset for CS-1, 2 and 3.

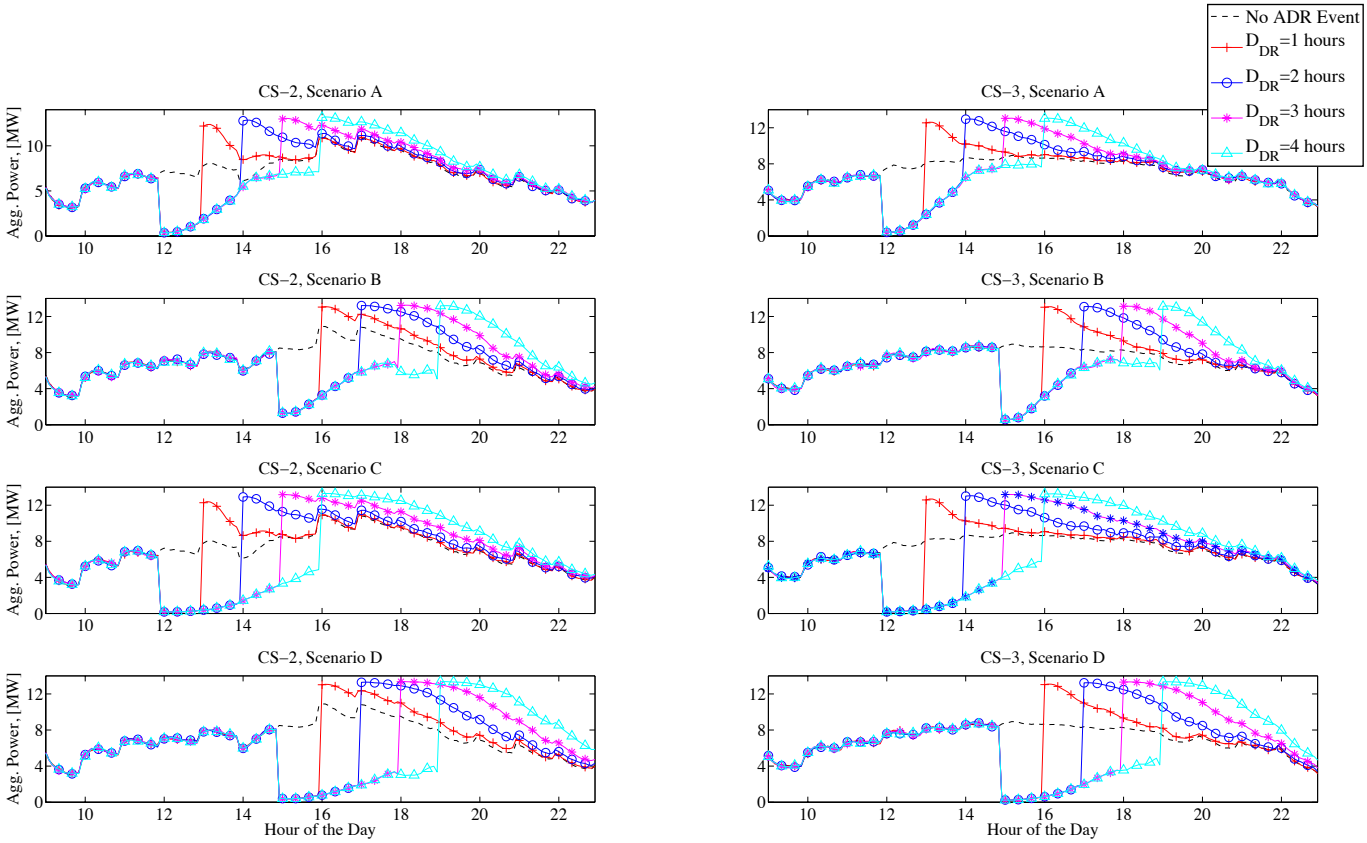

Figure 7. The behavior of the aggregate power consumption of the HVAC population for CS-2 and 3. The markers do not represent individual data points; they are placed to differentiate different DR curves from each other.

5 Results

Figure 6 shows the empirical cumulative distribution functions of $d_{i,t}^{shed}$ for the constant exogenous input case (CS-1), weekdays (CS-2) and weekends (CS-3). In all case studies, for any shed duration value d such that $d \leq 4$, it is possible to see that scenarios C and D have a lower $Pr(d_{i,t}^{shed} \leq d)$ value. This is expected because the setpoint change value for both C and D is 4^oF, hence given a shed duration value, there are less loads available provide the requested shed at least for d long. For CS-1, since the exogenous inputs are assumed to be constant during the day, scenarios A, B, C and D yield to almost identical results. When the results for CS-1 is compared to CS-2 and 3, it is possible to see that constant exogenous input assumption yields to a significant over-estimation of the number of loads that can provide a requested shed duration d for scenarios C and D. For scenarios A and B no significant difference is observed, which can be explained by the presence of less variation in the exogenous inputs during shorter shed periods that can be achieved by a lower setpoint change. For CS-2 and 3, it seems that scenario C has the lowest $Pr(d_{i,t}^{shed} \leq d)$ value among all 4 scenarios. This suggests that an ADR event at 12PM has the highest chance of exceeding a shed duration of 4 hours for both weekdays and weekends. Figure 7 shows the aggregate power time series obtained for each scenario under CS-2 and CS-3 against a no ADR event case with identical input conditions. We choose to depict CS-2 and 3, since they represent a more realistic scenario. For all the scenarios with an ADR event starting at $12PM$, it is possible to see that the recovery peak value is increasing with the increasing D_{DR}. Since the maximum value of D_{DR} we investigated in this study is 4 hours, the benefits due to decreasing exterior temperature (if any) seems to be diminished. In contrast, for the scenarios with ADR events scheduled at $3PM$, a decrease in the rate of increase of the aggregate power consumption is often visible after $6PM$, which helps to decrease the expected increase in the power consumption during recovery.

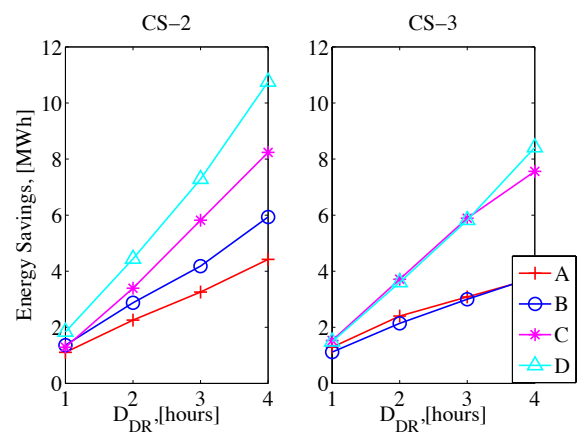

Figure 8. Energy savings w.r.t. no ADR case for CS-2 and 3 between 9AM and 11PM.

In Figures 8, 9, and 10, we depict the aggregate energy savings during the entire simulation period, energy savings during the event

Table 2. Characteristics of different scenarios and case studies

Case Study	External Temp.	Setpoint Profile	Setpoint Change $\Delta T^{set}_{t_{event}}$	Scenarios	Event Start Hour, t_{event}, [hour of the day]	DR event Duration, D_{DR}, [hours]
CS-1	82^oF, Constant	76^oF, Constant	2^oF	A	12	[1,2,3,4] hours
				B	15	
			4^oF	C	12	
				D	15	
CS-2	Measured, Thu, 21/06/2012	Avg. Weekday Profile	2^oF	A	12	[1,2,3,4] hours
				B	15	
			4^oF	C	12	
				D	15	
CS-3	Measured, Sun, 17/06/2012	Avg. Weekend Profile	2^oF	A	12	[1,2,3,4] hours
				B	15	
			4^oF	C	12	
				D	15	

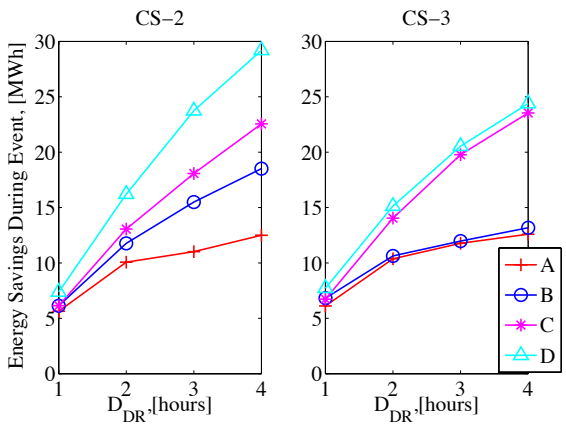

Figure 9. Energy savings during event w.r.t. no ADR case for CS-2 and 3.

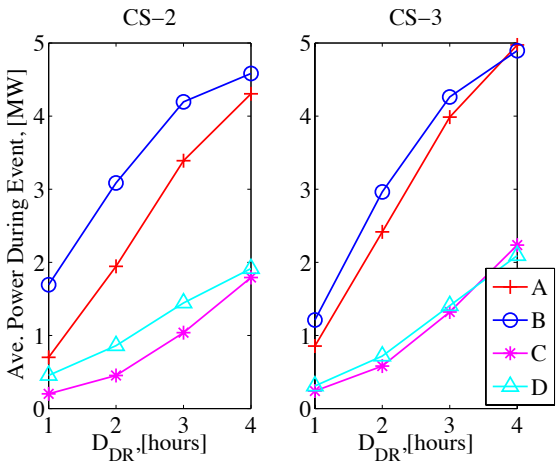

Figure 10. Average power during event for CS-2 and 3.

and the average power consumption during the ADR event. It is possible to see that scenarios D and C can offer more overall energy savings than scenarios A and B for weekdays and weekends. It is also possible to see that a later ADR event ($3PM$ in this case) can provide additional benefits during weekdays, which can be explained by the expected increase in the occupancy in households after earlier events resulting in a lower recovery setpoint. However, on the weekends, there is no significant difference between earlier and later events in terms of overall energy savings. In addition, the rate of increase in the energy savings with increasing D_{DR} shed duration during the ADR event diminishes with increasing D_{DR}. Even though some of the HVAC cannot provide a full shed during a long D_{DR}, their delayed schedule benefits highly from the decrease in the external temperatures.

6 Conclusions

In this paper, we have used a unique dataset to estimate and validate ARMA models that capture the thermal dynamics of individual HVAC units. Using these models, we investigated the potential of ADR that can be provided with an aggregation of HVACs under certain conditions. We analyzed the variability in duration of the shed each HVAC unit can provide and we quantified the energy savings that can be achieved through various setpoint adjustment scenarios.

In the individual ARMA models, the estimated solar heat gains are often counter intuitive. They are relatively small and equally likely to be positive as they are negative, signifying the implausible interpretation that solar insolation acts as a cooling source in 50% of residences. We believe that this is an effect of using a solar-earth geometry model that does not account for cloud cover, and thus poorly represents variation in insolation inside the home. Instead, the regular diurnal pattern may correlate with occupancy, where residents are less likely to be home during the high insolation time (around noon). This may explain the negative coefficients for some homes, as high solar times correlate with low occupancy, reducing internal heat gains. Another reason may be the correlation between the outside air temperature and the regular diurnal insolation pattern. The colinearity of insolation with occupancy and/or outside air temperature will be further investigated in future work.

Assuming a constant coefficient of performance of the air conditioner also affects our estimates. We expect that this assumption causes us to over-predict the power required to achieve a certain level of cooling when the indoor/outdoor temperature differential is low and to under-predict the required power when the differential is high. Thus, on a very hot day, we should expect that the instantaneous savings from an ADR event will be greater than predicted,

as the HVAC unit will be operating at a lower COP.

Our validation suggests that we are over-predicting the thermal capacity of the building, causing our model to predict longer than observed durations of transient events following a set point change. This over-prediction could be caused by rounding errors in our low resolution data. For instance, when the room temperature is maintained at a set point, it is often measured to be constant when in fact we know it to be fluctuating within a dead band. Thus, the effect of HVAC cycling on internal temperature is masked by rounding during these times.

As discussed in Section 5, the proposed demand response potential estimation strategy provided important insights into the response characteristics of an aggregate load population based on the data collected from individual households. In particular, we developed a proof-of-concept DR potential estimation strategy from a large aggregation of residential loads using ARMA models to represent individual thermal dynamics of HVAC units available in over 4200 households.

For future work, we propose to use an instrumental variable (changes in the set point) to predict HVAC energy use during times when the internal temperature is constant, thus removing the variation caused by cycling. We expect that this will increase the coefficients for HVAC unit power and decrease the estimated thermal capacity. We also would like to incorporate stochasticity in the DR potential estimation methodology used in this paper to reflect the confidence boundaries obtained from the individual ARMA models of HVACs.

7 Acknowledgments

Part of the work described in this paper was coordinated by the Consortium for Electric Reliability Technology Solutions and was funded by the Office of Electricity Delivery and Energy Reliability, Transmission Reliability Program of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

8 References

[1] H. Aalami, M. P. Moghaddam, and G. Yousefi. Demand response modeling considering interruptible/curtailable loads and capacity market programs. *Applied Energy*, 87(1):243–250, 2010.

[2] S. Bashash and H. Fathy. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In *American Control Conference (ACC), 2011*, pages 4546–4553, 2011.

[3] D. S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(5):1389–1400, May 2009.

[4] P. Cappers. Mass market demand response and variable generation integration issues: A scoping study. 2012.

[5] M. E. Dyson, S. D. Borgeson, M. D. Tabone, and D. S. Callaway. Using smart meter data to estimate demand response potential, with application to solar energy integration. *Energy Policy*, (0):–, 2014.

[6] J. H. Eto. Demand response spinning reserve demonstration–phase 2 findings from the summer of 2008. *Lawrence Berkeley National Laboratory*, 2010.

[7] A. H. R. Faruqi, G. S.S., J. Bode, P. Mangasarian, I. Rohmund, G. Wikler, and S. Glosh, D. Yoshida. A national assesment of demand response potential. Technical report, Federal Energy Regulatory Commission (FERC), June 2009.

[8] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Frequency regulation from flexible loads: Potential, economics, and implementation. 2013.

[9] M. Hummon, D. Palchak, P. Denholm, J. Jorgenson, D. J. Olsen, S. Kiliccote, N. Matson, M. Sohn, C. Rose, J. Dudley, et al. Grid integration of aggregated demand response, part 2: Modeling demand response in a production cost model. 2013.

[10] E. C. Kara, M. Berges, B. Krogh, and S. Kar. Using smart devices for system-level management and control in the smart grid: A reinforcement learning framework. *IEEE SmartGridComm'12*, Nov. 2012.

[11] S. Koch, J. L. Mathieu, and D. S. Callaway. Modelling and control of aggregated heterogenous thermostatically controlled loads for ancillary services. *Power Systems Computation Conference*, Aug. 2011.

[12] N. Lu, D. Chassin, and S. Widergren. Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model. *IEEE Transactions on Power Systems*, 20(2):725–733, May 2005.

[13] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld. Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. *Proceedings of the ACEEE Summer Study on Buildings, Pacific Grove, CA*, 2012.

[14] J. L. Mathieu. Modeling, analysis, and control of demand response resources. 2013.

[15] M. A. Piette, D. Watson, N. Motegi, and S. Kiliccote. Automated critical peak pricing field tests: 2006 pilot program description and results. *Lawrence Berkeley National Laboratory*, 2007.

[16] A. Rabl. Parameter estimation in buildings: methods for dynamic analysis of measured energy use. *Journal of Solar Energy Engineering*, 110(1):52–66, 1988.

[17] J. Short, D. Infield, and L. Freris. Stabilization of grid frequency through dynamic demand control. *Power Systems, IEEE Transactions on*, 22(3):1284 –1293, Aug. 2007.

[18] J. Taneja, K. Lutz, and D. Culler. The impact of flexible loads in increasingly renewable grids. In *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, pages 265–270. IEEE, 2013.

[19] K. Turitsyn, S. Backhaus, M. Ananyev, and M. Chertkov. Smart finite state devices: A modeling framework for demand response technologies. Mar. 2011.