

The value of real-time data in controlling electric loads for demand response

Johanna L. Mathieu, Duncan S. Callaway

Abstract—We describe approaches to controlling aggregations of residential thermostatically controlled loads (TCLs), such as air conditioners and refrigerators, to participate in energy and ancillary services markets, and energy arbitrage. Additionally, we examine the effect of offline and real-time data on the ability of the aggregations to track signals or minimize energy costs. We find that if real time information is available only from the distribution substation, load aggregations can track market signals, though tracking results improve significantly if each load transmits its temperature and on/off state in real time. Real-time information is also important for energy arbitrage. We find that with real-time data from the distribution substation and a good (offline) understanding of a TCL aggregation’s aggregate power and energy capacities as a function of ambient temperature, it is possible for TCL aggregations to save money on energy purchased in 5-minute energy markets. Improved access to real-time data (e.g., from the individual loads) and more advanced controls approaches may yield even better results. In sum, our results shed light on the required sensing, communications, and data needed for loads to effectively participate in electricity markets and arbitrage.

I. INTRODUCTION

Deployment of advanced metering infrastructure promises increasing access to data about residential end use electricity consumption. Real-time data could be used to achieve fast time-scale demand response (DR), allowing loads to participate in both energy and ancillary services markets. Importantly, we expect DR will be a low-cost resource capable of providing much of the additional regulation, load following, and ramping required of power systems with high penetrations of variable and uncertain renewables. However, how much and what kind of data do we really need to control loads with high fidelity? What is the value of offline versus real-time data? How do data needs change if loads participate in electricity markets or energy arbitrage?

Our work focuses on using offline and real-time data along with load aggregation models, state estimation techniques, and control and optimization strategies, to control heterogeneous populations of residential thermostatically controlled loads (TCLs) to deliver power systems services and participate in short time scale energy markets. TCLs operate within a hysteric dead-band and therefore have inherent energy storage capabilities. In this white paper, we summarize our results describing the role of real-time data in energy and ancillary




Fig. 1. A cooling TCL’s temperature dead-band and discretized state bins.

services market signal tracking and we describe our preliminary work to understand data and modeling requirements for energy arbitrage.

II. ELECTRICITY MARKET SIGNAL TRACKING

The goal of the central controller is to actuate TCLs to track a market dispatch or automatic generation control signal. We assume that the controller can switch loads on or off while the loads are within their temperature dead-band; however, it can not change a TCL’s temperature directly or affect its temperature set point. Additionally, we assume TCLs become uncontrollable if they are outside of the dead-band. This ensures that the TCLs provide the service requested by the user. The performance of the central controller is a function of the offline and real-time information available for system identification, state estimation, and control. In this section, we summarize our recent findings. More information is available in [1], [2], [3].

We modeled individual TCLs with difference equations and, for control and state estimation purposes, we developed an LTI model of populations of heterogenous TCLs:

$$x(k + 1) = Ax(k) + Bu(k) \tag{1}$$

$$y(k) = Cx(k) \tag{2}$$

where x contains the fraction of TCLs in discretized state bins characterizing both a TCL’s current temperature and on/off state (Fig. 1), A is the transpose of the Markov transition matrix, u is the input vector which allows us to switch TCLs in certain bins on or off, B ensures that TCLs maintain their temperature state while switching from on to off or vice versa, and y is the aggregate power consumption of the TCL population. This model is similar to those in [4], [5], [6].

J.L. Mathieu is with the Department of Mechanical Engineering at the University of California at Berkeley, 4th Floor Collaboratory, Sutardja Dai Hall, Berkeley, CA 94720-1740 USA, email: jmathieu@berkeley.edu.

D.S. Callaway is with the Energy and Resources Group and the Department of Mechanical Engineering at the University of California at Berkeley, 310 Barrows Hall, Berkeley, CA 94720-3050 USA, email: dcal@berkeley.edu.

TABLE I

CONTROLLER TRACKING PERFORMANCE: SEVERAL DATA SCENARIOS (RMS ERROR AS A PERCENT OF STEADY STATE POWER CONSUMPTION OF A 1,000 TCL POPULATION)

real-time data:	controller: offline data:	PPC (1) or (2)	PPC (3)	SPC (3)
(a) full state, each TCL		0.59%	x	1.1%
(b) partial state, each TCL		0.69%	x	1.1%
(c) partial state, 30% of TCLs		4.8%	x	4.3%
(d) substation, 5% forecast error		5.2%	7.1%	9.1%
(e) substation, 10% forecast error		6.1%	8.1%	10.5%

x: scenario unrealistic and not investigated

To minimize communication from the central controller to the TCLs, we broadcast a control vector of ‘switch probabilities’ to all TCLs, and TCLs decide whether or not to switch probabilistically based on their current temperature and on/off state, which corresponds a specific entry of the control vector. We examined three scenarios for offline data and five scenarios for real-time data sent from the TCLs to the central controller every 2 seconds.

Offline data scenarios:

- (1) TCLs report their physical parameters, ambient conditions, and set points (each assumed time invariant)
- (2) TCLs record and report their on/off state and temperature over some period of time
- (3) no data is available offline

Real-time data scenarios:

- (a) TCLs report their current temperature and on/off state (i.e. full state)
- (b) TCLs report their on/off state (i.e. partial state)
- (c) 30% of the TCLs report their on/off state (i.e. partial state)
- (d) the distribution substation reports aggregate power – this compared to the counterfactual substation load forecast gives us a noisy estimate of the aggregate power consumed by the TCLs – and the load forecast error is Gaussian white noise with standard deviation 5%
- (e) same as scenario d, except the load forecast error standard deviation is 10%

We designed two controllers: a predictive proportional controller (PPC) which uses all available offline and real-time data to identify the LTI model and develop state estimates with Kalman Filters in real time, and a simple proportional controller (SPC) which simply uses noisy measurements of y in real time. For the PPC, if offline data was available (scenario 1 or 2), we identified a 40-bin TCL model, and, if it was not, we identified a 2-bin TCL model using a joint parameter and state estimation technique. In simulation, we controlled populations of individual TCLs to track trajectories designed to mimic California ISO 5-minute market signals, increasing or decreasing the TCL population’s aggregate power consumption by up to one-quarter of the steady state aggregate power consumption. In each case, we tuned the controller gains to minimize RMS tracking error. Results are shown in Table I.

The PPC almost always performs better than the SPC meaning that there is value in the modeling and state estimation framework. The one exception is scenario c. In that case, the

measurement noise is not Gaussian so the Kalman Filter has trouble estimating the state. The results also show the value of offline and real-time data. If real-time state information from each TCL is available, the TCL population can follow the signal very closely. Interestingly, if real-time information is only available from the distribution substation, the TCLs can still track the signals though the tracking performance degrades. These results point out the important tradeoff between performance and costs.

III. ENERGY ARBITRAGE

Here, we examine the potential for energy arbitrage, i.e. consuming more energy when the wholesale electricity price is high and less when the price is low, in 5-minute energy markets. In this case, the TCL population is not sent a control signal to follow. Instead, TCL populations must compute their own trajectory, with the goal of minimizing their total energy costs. We employ a hierarchical approach consisting of three layers of control:

- 1) a linear program (LP) which solves for the optimal control trajectory based on a simple energy storage model
- 2) a TCL population controller (e.g., the PPC or the SPC described in the previous section) which controls TCLs to follow the optimal trajectory
- 3) a population of TCLs (each with a hysteric local controller), each modeled with a difference equation

This approach can be thought of as similar to approaches commonly used in aviation path planning in which simplified models are used to compute trajectories and more precise and complicated models are used for control [7].

A TCL’s energy usage is a function of ambient temperature, and while a refrigerator’s or water heater’s ambient temperature (indoor temperature) may not vary significantly, air conditioners and heaters experience a wide range of ambient temperatures (outdoor temperature). In the analysis described in the previous section, we ignored this issue because it would not change our methodology, though it would require occasional re-identification of the model used in the PPC formulation (scenario 1 or 2). However, we can not ignore changing ambient temperatures when computing optimal control trajectories.

Consider a population of air conditioners. The population’s energy capacity (i.e. state of charge), e , in each time step, k , is related to past mean power usages, p :

$$e_{k+1} = e_k + (p_k - p_{\text{baseline},k})\Delta T \quad (3)$$

where ΔT is the length of each time step and p_{baseline} is the power usage in the absence of control, which is a function of the current ambient temperature *and* the past ambient temperatures. For example, an air conditioner does not immediately turn on in the morning when the outdoor air temperature rises above the TCL dead-band because the house has stored ‘coolth’ from the night.

e and p are positive and bounded by maximum capacities:

$$0 \leq e_k \leq e_{\text{max},k} \quad (4)$$

$$0 \leq p_k \leq p_{\text{max},k} \quad (5)$$

Since, as before, we assume that we can not control TCLs outside of the dead-band, e_{\max} is a function of the dead-band. Specifically, e_{\max} is the energy required for all air conditioners that are currently available for control (i.e. TCLs operating within the dead-band) to traverse from right edge of the dead-band to the left edge of the dead-band (see Fig. 1). p_{\max} is the aggregate power consumption of the TCL population when all available air conditioners are turned on at the same time. Importantly, e_{\max} and p_{\max} are also functions of both the current ambient temperature and the past ambient temperatures.

For arbitrage, our goal is to find $p^* = p_1, \dots, p_N$, which minimizes total energy costs over N time steps:

$$J = l_1 p_1 \Delta T + l_2 p_2 \Delta T + \dots + l_N p_N \Delta T \quad (6)$$

subject to (3)-(5), where l_k is the energy price in each interval. This problem can be solved as a linear program (LP):

$$\min f'x \quad (7)$$

$$\text{s.t } Ax \leq b \quad (8)$$

$$\text{and } x \geq 0 \quad (9)$$

where $x = [p_1, \dots, p_N, e_1, \dots, e_N]'$ and the constraints are given in (3)-(5).

In our preliminary implementation, we computed p_{baseline} , e_{\max} , and p_{\max} as a function of only the current ambient temperature. Specifically, we assume TCLs are available for control if the current ambient temperature is above the TCL dead-band. Using 2010 interval locational marginal prices from California ISO's Merced node and outdoor air temperatures from NOAA, we implemented an iterative LP with noisy price and temperature forecasts. Specifically, using a prediction horizon of one hour, we re-solved the LP in each time-step with updated price and temperature forecasts, but we only implement the first control in each step. In the future, we plan to investigate more sophisticated approaches, e.g., [8].

We modeled the temperature forecast error in the next hour as Gaussian white noise with standard deviation β , and we modeled the price forecast error, ω , with an AR(1)-process:

$$\omega_{k+1} = \gamma \omega_k + \alpha \epsilon_k \quad (10)$$

where ϵ is Gaussian white noise with standard deviation 1, and $\omega(1) = 0$. Varying α , β , and γ , we can understand how performance changes with forecast error. In Fig. 2, we show an example optimal control trajectory for approximately one day assuming $\alpha = 40 \text{ \$/MWh}$, $\beta = 3^\circ \text{C}$, and $\gamma = 0.3$.

After solving for the trajectory, we controlled the population to follow it using the SPC from the previous section. Using the PPC would require a time-varying model (a subject for future research). While the results of the LP predict significant energy cost savings, the TCL population is unable to track the signal closely resulting in lower energy savings (Fig. 3). For example, considering the same error statistics used to generate Fig. 2, the LP predicts 15% annual energy costs savings while the TCL population is only able to achieve about 9% annual energy cost savings. Savings predicted by the LP and actual savings over a range of α are shown in Fig. 4. The discrepancy between the LP-predicted savings and




Fig. 2. Optimized state of charge and control trajectory, given 5-minute energy prices and 1,000 air conditioners.

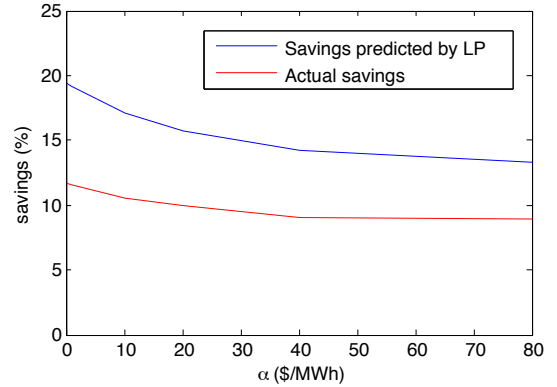


Fig. 4. LP-predicted and actual energy cost savings as a function of the standard deviation of the price forecast error, α . ($\beta = 3^\circ \text{C}$, $\gamma = 0.3$)

the actual savings is in part because we do not take into account temperature history when solving for p_{baseline} , e_{\max} , and p_{\max} causing us to significantly over estimate the resource in the mornings and somewhat underestimate the resource in the evenings. It is also because the proportional controller is unable to switch the correct number of TCLs in each time step, and this issue is further compounded by the fact that at each time interval (especially when the system is forced to the min or max) some TCLs are outside of the dead-band and therefore uncontrollable. It is likely that the PPC approach could lessen these issues.

For this analysis, we have assumed the controller has access to the aggregate power of the TCL population, and noisy price and temperature forecasts. Additional real-time information such as notification from each load when it becomes available for control, its current temperature, and its current on/off state would help us compute p_{baseline} , e_{\max} , and p_{\max} , though we have not yet determined the value of that information. Additionally, we plan to explore the benefit of including real-




Fig. 3. TCL population tracking performance: 1,000 air conditioners.

time information about the TCL population dynamics in the iterative LP. Right now the trajectory is created open-loop with respect to the TCL population dynamics (though closed-loop with respect to the price forecasts).

IV. CONCLUSIONS

TCLs could provide a significant amount of the energy storage capacity desired in power systems with high penetrations of intermittent renewables. However, there are many open questions about the best methods of control and the sensing, communications, and data required for each method. We have detailed some of the cost/performance trade-offs to consider when using TCLs in electricity markets, and energy arbitrage. Our results show that TCL aggregations are able to provide reliable services. Additionally, our preliminary research into energy arbitrage show that significant energy cost savings are possible.

V. ACKNOWLEDGEMENTS

We acknowledge Stephan Koch and Mark Dyson for related collaborative work and Maryam Kamgarpour for helpful insights. This work is supported in part by a PSERC Future Grid Initiative Grant.

REFERENCES

- [1] S. Koch, J. Mathieu, and D. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proceedings of the Power Systems Computation Conference*, Stockholm, Sweden, Aug. 2011.
- [2] J. Mathieu and D. Callaway, "State estimation and control of heterogeneous thermostatically controlled loads for load following," in *Proceedings of the Hawaii International Conference on Systems Science*, Wailea, HI, Jan. 2012.
- [3] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems (in review)*, 2012.
- [4] N. Lu, D. P. Chassin, and S. E. Widgren, "Modeling uncertainties in aggregated thermostatically controlled loads using a state queuing model," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 725–733, 2005.
- [5] N. Lu and D. Chassin, "A state queuing model of thermostatically controlled appliances," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1666–1673, 2004.
- [6] S. Bashash and H. K. Fathy, "Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads," in *Submitted to American Control Conference 2011*, 2011.
- [7] M. Kamgarpour, V. Dadok, and C. Tomlin, "Trajectory generation for aircraft subject to dynamic weather uncertainty," in *49th IEEE Conference on Decision and Control, Atlanta, GA*, 2010.
- [8] A. Bemporad, F. Borrelli, and M. Morari, "Model predictive control based on linear programming – the explicit solution," *IEEE Transactions on Automatic Control*, vol. 47, no. 12, pp. 1974–1985, 2002.