Benchmarking of Aggregate Residential Load Models Used for Demand Response

Gregory S. Ledva
Sarah Peterson
Johanna L. Mathieu
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, USA

***Abstract*—Residential loads can provide ancillary services such as frequency regulation to an electric power system. Aggregate load models aim to capture the dynamics of the demand-responsive loads in an accurate and computationally-tractable way, which allows the models to be incorporated into controllers and observers used by demand response providers. A variety of aggregate models have been developed; however, their accuracy has not been benchmarked against one another in comparable scenarios. This paper compares the accuracy of two Markov-based and one transfer function-based aggregate air conditioner (AC) models against a realistic simulation model with time-varying outdoor air temperature and temperature-dependent AC parameters. We also extend the existing models to cope with the time-varying outdoor air temperature. We find that 1) the more detailed Markov model is more accurate, 2) updating the Markov transitions as a function of the outdoor temperature trend decreases prediction error in both Markov models, 3) the transfer function model performs worst, likely because the simulation scenario differs significantly from the assumptions used to develop the model.**

***Index Terms*—Aggregate load model, demand response, electric power systems, frequency regulation, residential demand**

I. Introduction

The power consumption of large numbers of thermostatically controlled loads (TCLs), such as residential air conditioners (ACs), can be coordinated to help the electric power grid balance supply and demand [1], [2]. In addition to participating in traditional demand response programs, loads can be controlled to provide ancillary services, such as frequency regulation, by decreasing/increasing consumption with respect to their baseline. Much of the work on load control for ancillary services assumes a load aggregator receives a signal from the system operator and controls an aggregation of loads to match that signal. A significant body of recent work has sought to develop aggregate residential load models, e.g., [3]–[6], which the aggregator could use in state estimation and control algorithms. These models capture the dynamics of the total power consumption of the demand-responsive loads. Use of dynamic models generally improves control performance as compared to model-free control approaches. However, modeling aggregations of loads by representing each load individually leads to large and often complex models. For example, TCLs are best modeled as hybrid systems since they cycle on/off within a temperature hysteresis band. A good aggregate model balances simplicity and performance.

Despite significant recent efforts to develop aggregate load models, we only have a partial understanding of which models work best under which conditions. Each model was developed under a different set of assumptions (e.g., homogeneous vs. heterogeneous loads, static vs. dynamic ambient conditions such as outdoor temperature) and assuming a specific type of control (e.g., on/off switching, temperature setpoint control). Models are generally validated in simulation studies, and, often, the same unrealistic assumptions used to build the model are used to validate it.

In this paper, we seek a better understanding of the advantages and disadvantages of three aggregate load models that represent the dynamics of a heterogeneous AC population. Two of the models use Markov chains [3], [4] and one uses a transfer function [5], [6]. We identify each model and then use it to predict the aggregate power consumption of 10,000 air conditioners over a 24 hour period with a time-varying outdoor temperature. Prediction accuracy is computed against a realistic simulation model in which individual air conditioners are represented with hybrid system models including three states: indoor air temperature, indoor mass temperature, and on/off mode. We assume each AC has a time-varying cooling capacity, a time-varying coefficient of performance (COP), and a time-varying power draw (when the AC is on) that all depend on the outdoor temperature, making our plant more accurate than that typically used in the aggregate load modeling literature. In this preliminary work, we assume that the loads are not coordinated by a load aggregator; investigating model performance under aggregator control is a subject for future work.

Beyond the simulation-based benchmarking of three aggregate load models, this paper also includes several methodological contributions required to enable a fair comparison in a realistic setting with a time-varying outdoor air temperature. In particular, we extend the Markov models [3], [4] to update the Markov transition matrix parameters as a function of the temperature trend, and we discuss other ways in which the model could be improved. We also determine a transfer function structure that may lead to improved model accuracy as compared to the model in [5], [6] and provide suggestions on better ways to identify its parameters.

This research was funded by NSF Grant #ECCS-1508943.

TABLE I
AIR CONDITIONER MODEL PARAMETERS

Parameter	Description	Value
θ^{set}	Temperature Setpoint [°C]	$\mathcal{U}\{20, 24\}$
θ^{db}	Temperature Deadband [°C]	$\mathcal{U}\{1, 2\}$
θ_t^{o}	Outdoor Temperature [°C]	Varies from 25.2 to 35.0
U^{m}	Mass Conductance $[\frac{\text{kW}}{^\circ\text{C}}]$	$\mathcal{U}\{4.4, 5.4\}$
U^{a}	Air Conductance $[\frac{\text{kW}}{^\circ\text{C}}]$	$\mathcal{U}\{0.25, 0.3\}$
Λ^{m}	Mass Heat Capacitance $[\frac{\text{kWh}}{^\circ\text{C}}]$	$\mathcal{U}\{2.0, 2.5\}$
Λ^{a}	Air Heat Capacitance $[\frac{\text{kWh}}{^\circ\text{C}}]$	$\mathcal{U}\{0.5, 0.6\}$
Q_t^{AC}	Cooling Capacity [kW]	$\frac{Q^{\text{rate}}\ (1.32-0.01\ \theta_t^{\text{o}})}{1+f^{\text{latent}}}$
Q^{rate}	Rated Cooling Capacity [kW]	$\mathcal{U}\{11.1, 13.5\}$
f^{latent}	Fraction of Latent Cooling [-]	0.35
P_t^{AC}	Power Draw [kW]	Q_t^{AC}/η_t
η_t	Coefficient of Performance [-]	$\eta^{\text{std}}(0.33 + 0.02\ \theta_t^{\text{o}})^{-1}$
η^{std}	COP at Standard Conditions [-]	3.5

The rest of this paper is organized as follows: Section II details the individual TCL model used in the plant and Section III describes each aggregate load model and our extensions. Section IV describes the simulation setting and results. Section V concludes.

II. INDIVIDUAL TCL MODEL

We use the hybrid model from the residential module of GridLAB-D [7] to model an individual AC within the plant. The hybrid model contains continuous states corresponding to the internal air temperature and mass temperature $\theta_t^a, \theta_t^m \in \mathbb{R}$ and a discrete state corresponding to the on/off mode $m_t \in \{0, 1\}$, where the AC is drawing power if $m_t = 1$. Table I summarizes other parameters used within the model where $\mathcal{U}\{a, b\}$ indicates a uniform distribution between a and b. Table I includes time-varying values for the cooling capacity Q_t^{AC}, COP η_t, and power draw P_t^{AC}, which are all functions of the time-varying outdoor temperature θ_t^{o}. Note that we have neglected several aspects that are included within the GridLAB-D model: minimum cycling times, the power draw and heat injection from the circulation fan, and heating from solar irradiance and internal heat gains; including these are a subject for future work.

The update equations for the hybrid model are

$$\theta_{t+1} = A\ \theta_t + B\ (Q_t^{\text{AC}} m_t) + E\ \theta_t^{\text{o}}, \tag{1a}$$

$$m_{t+1} = \begin{cases} 0 & \text{if } \theta_{t+1}^a < \theta^{\text{set}} - \theta^{\text{db}}/2 \\ 1 & \text{if } \theta_{t+1}^a > \theta^{\text{set}} - \theta^{\text{db}}/2 \\ m_t & \text{otherwise,} \end{cases} \tag{1b}$$

where (1a) updates the temperatures with $\theta_t = \begin{bmatrix}\theta_t^a & \theta_t^m\end{bmatrix}^{\text{T}}$, and (1b) updates the on/off mode if the air temperature reaches the edge of the allowable temperature range. Matrices A, B, and E are computed by first constructing continuous-time matrices using the thermal parameters U^{m}, U^{a}, Λ^{m}, and Λ^{a}, and then discretizing the continuous-time matrices using the time-step $\Delta t = 2$ seconds. To simulate a set of n^{AC} ACs, we parameterize the set of ACs by independently selecting the relevant parameters from the distributions in Table I, where $\mathcal{U}\{\cdot, \cdot\}$ refers to a uniform distribution. We update each AC's temperatures and on/off mode by applying (1a) and (1b) with the AC's parameters.

III. AGGREGATE TCL MODELS

If n^{AC} is large (e.g., on the order of thousands), then incorporating the n^{AC} hybrid models into control and estimation algorithms can be computationally prohibitive. As a result, these algorithms often employ aggregate models, which model the behavior of the n^{AC} ACs using a single, simpler model. Here, we compare three aggregate models. The first model, developed in [3] and referred to as the two-state Markov model, defines a set of discrete bins based on θ_t^a and m_t, constructs an aggregate state as the portion of air conditioners in each bin, and it uses a Markov transition matrix to update the aggregate state of the model. The second model, developed in [4] and referred to as the three-state Markov model, is similar but defines a set of discrete state bins based on θ_t^a, θ_t^m, and m_t. The third model, developed in [5], [6], referred to as the transfer function model, maps changes in the outdoor temperature to changes in the steady state aggregate demand.

A. Two-State Markov Model

The two-state Markov model [3] uses an aggregate state $x_t \in \mathbb{R}^{2n^{\text{a}}}$, which is the portion of air conditioners in each of $2n^{\text{a}}$ discrete state bins. The discrete state bins are formed by dividing a normalized temperature deadband into n^{a} temperature intervals, and then creating two discrete states in each interval, one for air conditioners that are on and one for those that are off. An AC maps to a bin based on its air temperature and on/off mode. The two-state autonomous Markov model is

$$x_{t+1} = A_t\ x_t, \tag{2}$$

$$y_t = C_t\ x_t, \tag{3}$$

where $y_t \in \mathbb{R}^1$ is the aggregate demand, A_t is a transposed Markov transition matrix where the entries correspond to the probability of bin transitions within the time-step, and $C_t = n^{\text{AC}}\ \overline{P}_t^{\text{on}}\ \begin{bmatrix}0 & \dots & 0 & 1 & \dots 1\end{bmatrix}$ where the scalar $\overline{P}_t^{\text{on}}$ is the average power draw of air conditioners that are on. In [3] $A_t = A$, corresponding to a constant outdoor temperature. In [8], A_t is a function of the time-varying outdoor temperature.

Here, to compute A_t and $\overline{P}_t^{\text{on}}$, we first compute a set of time-invariant matrices $\mathcal{A}$ and average power draws $\mathcal{P}^{\text{on}}$ at different outdoor temperatures $\mathcal{T}^{\text{o}}$. Each element of $\mathcal{A}$ is computed by counting the bin transitions at the corresponding outdoor temperature from $\mathcal{T}^{\text{o}}$, and then normalizing the columns to sum to one. We then compute A_t by linearly interpolating the elements of the two time-invariant matrices corresponding to the temperature above and below θ_t^{o}. Each element of $\mathcal{P}^{\text{on}}$ is computed by calculating the average power draw of ACs that are on when θ_t^{o} is at the temperature from $\mathcal{T}^{\text{o}}$. We interpolate $\overline{P}_t^{\text{on}}$ in a similar manner.

We use three approaches to identify $\mathcal{A}$ and $\mathcal{P}^{\text{on}}$.

- MM2-C: We simulate the n^{AC} ACs for each integer temperature within $\mathcal{T}^{\text{o}}$, holding the temperature constant during each simulation. Each entry of $\mathcal{A}$ and $\mathcal{P}^{\text{on}}$ is computed with data from one simulation.
- MM2-V: We simulate the ACs with a time-varying outdoor temperature using historical temperature data. Each

entry of $\mathcal{A}$ and $\mathcal{P}^{\text{on}}$ is computed with data corresponding to outdoor temperatures nearest to the entry's associated temperature.

- MM2-S: In addition to simulating the ACs with a time-varying outdoor temperature, we create two sets $\mathcal{A}$ and $\mathcal{P}^{\text{on}}$, one for when θ_t^{o} is increasing and one for when θ_t^{o} is decreasing. This is justified because the interaction between θ_t^{a} and θ_t^{m} is different for an increasing versus decreasing θ_t^{o}.

B. Three-State Markov Model

The three-state Markov model [4] creates discrete bins, similar to those of the two-state Markov model, but based on both the air and mass temperatures. The deadband is divided into n^{a} temperature intervals for the air temperature, as in the two-state Markov model, and the deadband is divided into n^{m} mass intervals for the mass temperature. The aggregate state is then $x_t \in \mathbb{R}^{2n^{\text{a}}n^{\text{m}}}$, which is the portion of ACs in each of the $2(n^{\text{a}}n^{\text{m}})$ discrete state bins. The three-state autonomous Markov model is (2) and (3). As with the two-state Markov model, we define three methods of constructing A_t and $\overline{P}_t^{\text{on}}$: MM3-C, MM3-V, and MM3-S. They each construct the sets $\mathcal{A}$ and $\mathcal{P}^{\text{on}}$ using the same methodology as with the two-state Markov model. Note that this model is constructed assuming that measurements of the thermal mass are available within each residence, which are not available in practice.

Also, note that the structure (i.e., the location of zero and non-zero entries) of each time-invariant matrix within $\mathcal{A}$ can be different for this model, depending on the outdoor temperature used to construct the model. As A_t is computed by linearly interpolating between different time-invariant matrices with different structures, the structure of A_t changes over time. Non-zero elements of x_t can then be set to zero due to the changing structure, reducing the probability mass within x_t to less than one. We heuristically ensure that each column in the time-invariant matrices sum to one, ensuring that the probability mass within x_t is not reduced.

C. Transfer Function Model

The transfer function model, developed in developed in [5], [6], maps a change in the ambient temperature to a change in the steady state aggregate demand, and it has the form

$$y(t) = \Big(\sum_{i=1}^{n^{\text{AC}}} P_t^{\text{AC},i} \Big) \, \mathcal{L}\left\{ G(s) \right\}^{-1} \left(\theta_t^{\text{o}} - \theta^{\text{init}} \right) \tag{4}$$

$$G(s) = \frac{b_1 s + b_2}{s^2 + \zeta \omega_{\text{n}} \, s + \omega_{\text{n}}^2} \tag{5}$$

where $\mathcal{L}\{\cdot\}^{-1}$ is the inverse Laplace transform, θ^{init} is the initial outdoor temperature, s is the variable in the Laplace domain, ζ is the damping coefficient of an under-damped system, ω_{n} is the natural frequency of an under-damped system, b_1 and b_2 are coefficients, and $\sum_{i=1}^{n^{\text{AC}}} P_t^{\text{AC},i}$ is the total power draw if all ACs were on at time t. The coefficients and parameters in (5) are computed based on an assumed response to a step change in the outdoor temperature from θ^{init} to θ^{final}, where these temperatures must be assumed *a priori* to compute the model. The details about the calculation of (5) can be found in [5], [6], but some general aspects of the calculations include the following: 1) the assumed, individual AC model is a simplified version of that in Section II, which only contains one continuous state for the air temperature; 2) many of the parameters characterizing the individual AC model are assumed to be identical across the AC population; 3) estimates for the steady state portion of ACs that are on must be computed for both θ^{init} and θ^{final}; and 4) an estimate for the the period of the oscillations resulting from the step change is computed using the parameter distributions of the population, which are assumed to be known.

The AC models in Section II differ substantially from the assumptions within [6]. As a result, we define two approaches to generate models based on transfer functions.

- TF-O: We implement (4) and (5) according to [6]. Thermal parameters for the AC model with one continuous state are identified from historical data of θ_t^{a}, m_t, and θ_t^{o} data for each AC. Parameters that are assumed to be identical across the AC population in [6] are taken as their average value.
- TF-ID: We use a data-driven approach to identify a transfer function structure as well as the transfer function parameters from historical θ_t^{o} and aggregate demand data. We identify the transfer function parameters using the *tfest*$(\cdot)$ function in Matlab, where we subtract the initial values of the aggregate demand data and θ_t^{o} data. We then select the transfer function structure that achieves the lowest error based on historical data. The resulting model has the following structure:

$$y(t) = \mathcal{L}\left\{ \frac{b_3 s^2 + b_4 s + b_5}{s^2 + a_1 \, s + a_2} \right\}^{-1} \left(\theta_t^{\text{o}} - \theta^{\text{init}} \right) + y^{\text{init}} \tag{6}$$

  where y^{init} is the initial demand and the parameters to be identified are b_3, b_4, b_5, a_1, and a_2.

To implement each of the transfer functions within a discrete-time simulation, we convert each transfer function to a separate continuous-time, state-space model and then discretize each state-space model.

IV. Simulation-Based Case Studies

In this section, we use a number of simulation-based case studies to investigate the prediction accuracy of the three aggregate models and their variations. Section IV-A summarizes the simulation settings for the case studies, and Section IV-B presents the results. In general, more accurate aggregate models lead to more accurate control and state estimation algorithms that could be used to provide frequency regulation with ACs throughout the day. As a result, we evaluate the models using the RMS error (RMSE) of the model's predicted aggregate demand versus the true aggregate demand over the day.

A. Parameterization

The case studies each use a plant of 10,000 simulated ACs, modeled and parameterized according to Section II, over 24 hours using a time-varying outdoor temperature. The outdoor

TABLE II
SIMULATION SETTINGS

Parameter	Description	Value
Δt	Time-step duration [s]	2
n^{steps}	Number of time-steps [-]	43,200
n^{AC}	Number of air conditioners [-]	10,000
n^{a}	Number of air temperature bins [-]	20
n^{m}	Number of mass temperature bins [-]	20
$\mathcal{T}^{\text{o}}$	Set of temperatures used to compute the various Markov models	$\{24, 25, \ldots, 36\}$
θ^{init}	Initial outdoor temperature used to compute the transfer function [°C]	28.8
θ^{final}	Final outdoor temperature used when computing the transfer function [°C]	35.0

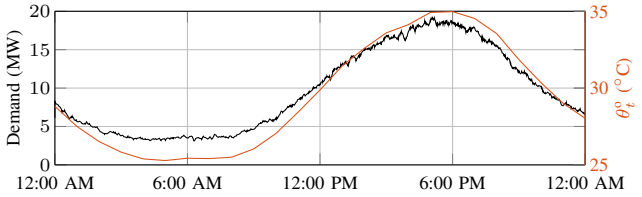

Fig. 1. The aggregate demand of the simulated ACs and the outdoor temperature for the case study setup

temperature, available through the Pecan Street Dataport [9], corresponds to that of Austin, TX on July 10, 2016, where we linearly interpolate the data from one-hour to two-second time-steps. Figure 1 depicts the aggregate demand of the simulated ACs, which is the ground-truth demand in our simulations, along with θ_t^{o}. Table II contains additional simulation parameters. We use the RMSE to evaluate the prediction accuracy of each aggregate model versus the aggregate demand of the simulated ACs.

To compute the Markov models, we generate the necessary data by simulating the plant models using outdoor temperature data from July 1-9, 2016. We initialize the Markov model at the true aggregate state value. To generate the parameters for the transfer function models, we simulate the plant models using the outdoor temperature for July 9, 2016. For TF-O, the transfer function was computed using θ^{init} as the actual initial temperature of the simulation and θ^{final} as the maximum temperature of the simulation. The initial state of the discreteized TF-O model is set to a vector of zeros. For TF-ID, y^{init} is set to the initial aggregate demand, θ^{init} is set to the actual initial temperature, and the initial state of the discretized TF-ID model is set to zeros.

B. Results

Table III summarizes the aggregate model variations and their RMSE values. Figure 2 presents time series of the aggregate demand and the predictions of a number of aggregate models over a portion of the simulated day. Figure 2a presents time series for MM2-C and MM2-S where we exclude MM2-V for clarity, and Fig. 2b presents time series for TF-O and TF-ID. We do not present time series for the three-state Markov model predictions as they are similar to those of the two-state Markov model. Figure 3 presents errors of various aggregate model's predictions versus the aggregate demand over the entire simulated day. Figure 3a presents the errors of MM2-C and MM2-S, and Fig. 3b presents the prediction error for

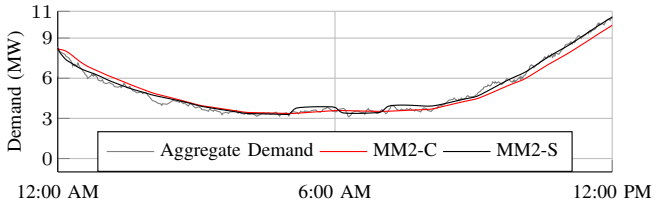

(a) Two-state Markov models

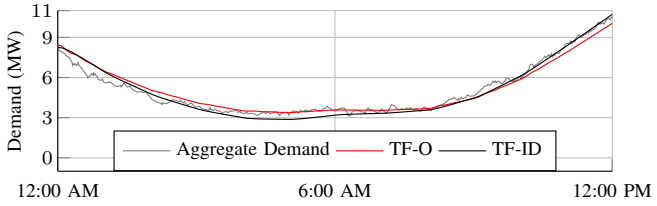

(b) Transfer function models

Fig. 2. Time series of the aggregate demand and various model predictions over a portion of the simulated day

MM3-C and MM3-S, where we exclude MM2-V and MM3-V for clarity. Figure 3c presents the prediction error for TF-O and TF-ID.

As can be seen in Table III, the increased modeling detail of the three-state aggregate model reduces the RMSE errors versus the two-state aggregate model across all variations. Accounting for the trend in θ_t^{o} in MM2-S and MM3-S reduces the RMSE versus the other Markov-based model varieties substantially. Furthermore, both MM2-S and MM3-S achieve similar RMSE, indicating that accounting for the θ_t^{o} reduces the gap in modeling accuracy between the two Markov-based aggregate models. The reduced RMSE when using MM2-S or MM3-S can be see in Fig. 2a (which shows MM2-S but is true for MM3-S as well). Whereas the prediction of MM2-C lags the aggregate demand, this is corrected in MM2-S by accounting for differences in the AC model behavior when θ_t^{o} is increasing versus decreasing. Also, note that from 4:00 AM to 8:00 AM, θ_t^{o} is relatively flat and switches between increasing and decreasing several times, and the resulting prediction of MM2-S undergoes several jumps. This results in significant prediction error in MM2-S and MM3-S, as can be seen in Fig. 3a and Fig. 3b. Further differentiating between an increasing, decreasing, or flat trend in θ_t^{o} may improve model performance. The time series for MM3-C and MM3-S, which are not included, show similar trends.

The predictions of TF-O and TF-ID have higher RMSE than either Markov-based aggregate model. This makes sense as the transfer function models are simpler models than the Markov-based models; whereas TF-O and TF-ID have two states and a steady-state demand term, the Markov-based models have 40 and 800 bins, respectively. Figure 2b shows that the TF-O prediction generally lags the aggregate demand and the TF-ID prediction does not. The errors in Fig. 3c also show that TF-O under-predicts the demand as it increases from 12:00 PM to 6:00 PM.

V. CONCLUSIONS

In this work, we benchmarked the prediction accuracy of three existing aggregate models and several variations of them

TABLE III
SUMMARY OF MODELS AND RMSE (KW) VALUES

Abbreviation	Base Aggregate Model	Details	RMSE (kW)
MM2-C	Two-State Markov Model	Set of models for different θ_t^o values; data for model computation generated using constant θ_t^o values	436.7
MM2-V	Two-State Markov Model	Set of models for different θ_t^o values; data for model computation generated using time-varying θ_t^o values	437.1
MM2-S	Two-State Markov Model	Set of models for different θ_t^o values and different θ_t^o trends; data for model computation generated using time-varying θ_t^o	226.2
MM3-C	Three-State Markov Model	Set of models for different θ_t^o values; data for model computation generated using constant θ_t^o values	320.9
MM3-V	Three-State Markov Model	Set of models for different θ_t^o values; data for model computation generated using time-varying θ_t^o values	322.9
MM3-S	Three-State Markov Model	Set of models for different θ_t^o values and different θ_t^o trends; data for model computation generated using time-varying θ_t^o	213.4
TF-O	Transfer Function Model	Single model; assumed transfer function structure of two poles and one zero; parameters computed using [6]	504.4
TF-ID	Transfer Function Model	Single model; transfer function structure of two poles and two zeros identified from historical model accuracy; parameters identified with historical input-output data	447.0

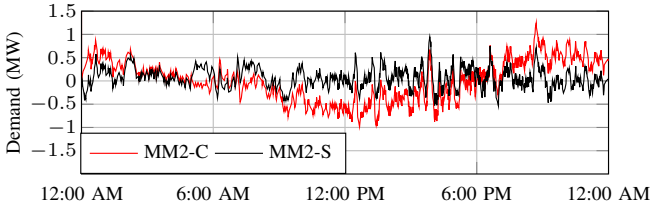

(a) Two-state Markov models

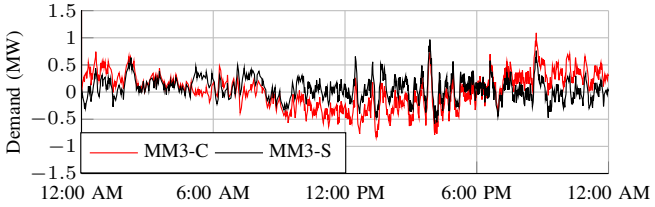

(b) Three-state Markov models

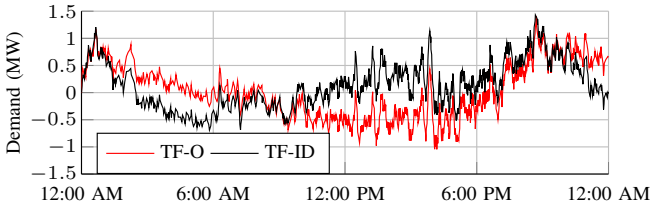

(c) Transfer function models

Fig. 3. Prediction error of various models versus the aggregate demand

within a common, detailed simulation scenario. The simulation scenario includes a time-varying outdoor temperature and detailed AC models that include a time-varying cooling capacity, COP, and power draw (when the AC is on) that all depend on the outdoor temperature. Results indicate that the three-state aggregate model generally performs better than the other aggregate models. Incorporating the temperature trend improves both Markov-based models and reduces the gap in prediction accuracy between the two. The transfer function model is the least accurate of the three aggregate models, most likely due to its simplicity and due to the fact that the simulation scenario differs substantially from the assumptions used to develop the model in [6]. The transfer function using input-output data to identify the parameters resulted in reduced prediction error. While the three-state aggregate model is the most accurate, it is more computationally complex than the other two models. The simpler two-state aggregate model offers similar performance, when including temperature trends into the model, at lower computational complexity.

Avenues of future work include the following: 1) developing a variation of the Markov-based models that accounts for times of little change in outdoor temperature; 2) deriving transfer function parameters for an AC aggregation undergoing a sinusoidal input rather than a step input, which better approximates realistic temperature changes; 3) investigating the identified transfer function structure and whether its parameters can be derived from the individual AC population; 4) including time-varying solar irradiance and internal heat gains within the AC models; and 5) prediction performance under aggregator control.

REFERENCES

[1] D. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, pp. 1389–1400, 2009.

[2] D. Callaway and I. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.

[3] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

[4] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[5] C. Perfumo, E. Kofman, J. H. Braslavsky, and J. K. Ward, "Load management: Model-based control of aggregate power for populations of thermostatically controlled loads," *Energy Conversion and Management*, vol. 55, pp. 36–48, 2012.

[6] N. Mahdavi, J. H. Braslavsky, and C. Perfumo, "Mapping the effect of ambient temperature on the power demand of populations of air conditioners," *IEEE Transactions on Smart Grid*, 2016.

[7] "GridLAB-D Residential module user's guide," http://gridlab-d.shoutwiki.com/wiki/Residential_module_user%27s_guide, Accessed: 2017-11-01.

[8] J. L. Mathieu, M. Kamgarpour, J. Lygeros, G. Andersson, and D. S. Callaway, "Arbitraging intraday wholesale energy market prices with aggregations of thermostatic loads," *IEEE Transactions on Power Systems*, vol. 30, no. 2, pp. 763–772, 2015.

[9] Pecan Street Inc., "Dataport," 2016.