Meters to Models:
Using Smart Meter Data to Predict Home Energy Use

Krystian X. Perez, Wesley J. Cole, Michael Baldea and Thomas F. Edgar
University of Texas at Austin

ABSTRACT

Access to smart meter data in the United States presents an opportunity to better understand residential energy consumption and energy-related behaviors. Air-conditioning (A/C) use, in particular, is a highly variable and significant contributor to residential energy demand. Most current building simulation software tools require intricate detail and training to accurately model A/C use within an actual house. However, integrating existing modeling software and empirical data has the potential to create highly portable and accurate models. Reduced-order models (ROM) are low-dimensional approximations of more complex models that use only the most impactful variables. In this paper, we report on the development of ROMs for 41 physical houses in Austin, Texas, using smart meter data. These models require outdoor dry bulb temperature, thermostat set points and A/C energy use data to regress model coefficients. A non-intrusive load monitoring technique is used to disaggregate A/C electricity consumption from whole-house electricity data reported by smart meters. Thermostat set points are provided by smart thermostats. Once trained, the models can use thermostat set points and dry bulb temperatures to predict A/C loads. The ROMs are used to evaluate the potential of automated thermostat control to reduce the aggregate peak demand. A centralized model predictive controller reduces the aggregate peak load by adjusting the thermostat set points to pre-cool houses and staggers the time A/C units turn on.

Introduction

Residential energy use is a significant contributor to energy demand during peak hours. In particular, A/C use is highly variable and an important contributor to fluctuations in demand. In ERCOT (the Electric Reliability Council of Texas), residential loads account for nearly 50% of the summer peak demand, driven primarily by A/C use (Wattles 2011). The magnitude of this demand is heavily influenced by ambient conditions, such as outdoor temperature, solar radiation and human behavior. Shaping the residential A/C load provides an opportunity to diminish reliance on inefficient peaking power plants. One method of leveling peak loads is to use energy models to identify the human activity and weather factors that impact energy consumption, predict energy use, and finally, optimize energy use through shifting cooling loads via an appropriate operating schedule for the A/C.

Residential energy modeling has been in practice in the United States for a number of years. Modeling software, such as eQuest and EnergyPlus work in a semi-empirical framework that utilizes first-principles equations to describe heat flow as well as empirical efficiencies of equipment in buildings (U.S. DOE 2014). The accuracy of the models has been evaluated using real energy data, particularly in commercial buildings, and results have varied (Lomas et al. 1997; Christensen et al. 2010). While commercial software products are able to represent a thermodynamic model of a house, they require many details to accurately describe the effect of behaviors and materials inherent in an actual house, such as insulation thickness and type, occupancy schedules and duct leakage. Also, the required training and time to calibrate the model is intensive and a deterrent to potential users in the residential sector where building types vary widely. While it may be reasonable to model a single house accurately, it would not be practical to accurately represent a group of houses or neighborhood. A recent observation on A/C behavior and energy consumption between houses revealed that the variance in consumption behaviors increases linearly with outdoor temperature (Perez et al. 2014). In other words, at high temperatures, when peak energy would be important, the discrepancy in energy use among houses is largest.

In place of complex house models, we propose to use smart meter data to create simple, accurate and data-driven models of the A/C energy consumption. Smart meter data inherently contain information on the physical characteristics of houses and on human behavior. This research seeks to develop a method to regress parameters of reduced order models (ROMs) using these newly available smart meter data. ROMs are low-dimensional approximations of more complex models that reduce the number of variables to those that have the most impact. A ROM is effective in modeling individual houses because the dynamics, time constants and thermal capacity of houses, can still be captured. Several model-reduction techniques have been applied to building models, such as the lumped capacitance, resistance-capacitance (RC), autoregressive with exogenous inputs (ARX) and other simplified models (see Gouda, Danaher, and Underwood 2002; Karmacharya et al. 2012; Malisani et al. 2010; Anthony 2011). Reduced models are used in controls and optimization to reduce computational complexity when real-time solutions are needed. Cole *et al.* (2013) developed a ROM from an EnergyPlus model that determined air conditioning electricity consumption based on dry bulb temperature (DBT), thermostat set point temperature (T_{sp}), and time of day.

In this paper, we report on the development of ROMs for individual houses using smart meter data. These models require ambient temperature, thermostat set points and A/C energy use (provided by smart meter data) to regress model coefficients. A non-intrusive load monitoring technique is used in order to separate A/C energy consumption from whole-house energy data reported by smart meters (Perez et al. 2014). From the regressed coefficients, the model is then able to use thermostat set points and dry bulb temperatures to predict A/C loads. Using these ROMs, automated thermostat control is investigated to determine electricity peak reduction.

Methodology

Method Outline

An overview of the model creation process is shown in Figure 1. Whole-house electricity consumption data, thermostat set points, and dry bulb temperatures for 41 physical houses in Austin, Texas, were provided by the Pecan Street Smart Grid Demonstration Project in 1-minute time intervals from to June 1st through July 20th, 2013 (details about the data are provided in the "Data" section). The ROM used in this investigation requires A/C energy usage, so A/C energy use was disaggregated from the overall energy profile using a technique previously developed in our research group and described in (Perez et al. 2014). In this technique, the magnitude of change in load that signals the A/C turning on or off is found, which is then used to identify on and off events of A/C use. The specific form and parameters of the ROM model are discussed in detail in the "Model Form" section. Generally, the ROM uses previous values of A/C energy, dry bulb temperature and indoor set point temperature to predict A/C energy use at the next hour.

The smart meter data were separated into a training dataset and a validation dataset. The training dataset contained approximately 75% of the smart meter data and was used to regress model coefficients. The validation dataset contained the remaining 25% of the data and were used for validation purposes.

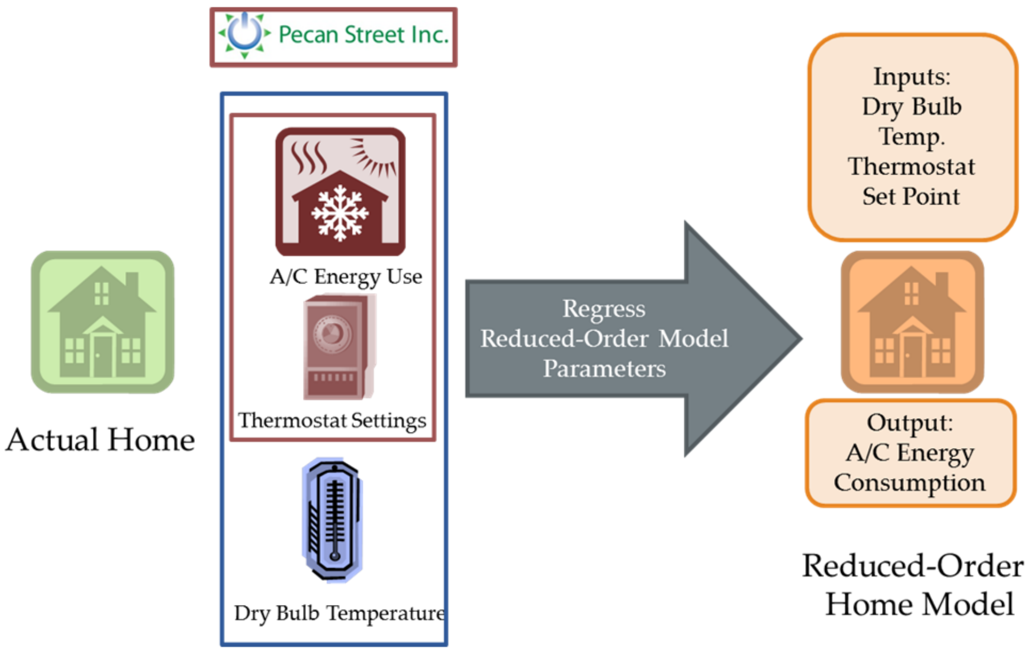

Figure 1. Diagram of individual house model creation process.

Data

Overall energy use and thermostat set points for 41 individual houses in the Mueller district, Austin TX were provided by Pecan Street Inc. in 1-minute time intervals. This granularity of data was chosen to provide accurate estimations of A/C load. However, previous work has shown that 5-minutes interval data is sufficient to separate A/C energy use from whole house energy data (Perez et al. 2014). Each house was metered with an eGauge power monitor that reported whole-house power consumption in watts (Rhodes et al. 2014). The homes were equipped with smart thermostats as well as that recorded the thermostat set point. The time period for data collection was limited to June 1st through July 20th, 2013 because smart thermostats that report set points have only been installed recently and further data were not available. In the ROM proposed by Cole *et al.* 2013, the model estimated A/C usage in hourly intervals using hourly inputs. The ROM hourly model was chosen because the model is linear, which simplifies the control and optimization problem later discussed. The on/off behavior of the A/C unit makes it difficult to fit an empirical model at finer time intervals and further research is required to develop an accurate model. Thus, the average A/C usage and temperature set points were averaged in one-hour time intervals. Outdoor air temperature was taken from a local weather station.

The Mueller neighborhood consists mostly of newer (since 2007), green-built houses and has a large amount of new technology penetration, such as rooftop photovoltaic panels and plug-in vehicles (Rhodes et al. 2014). The houses are equipped with electric A/C cooling units and natural gas heating systems. A/C energy will refer to electricity use that is used to cool the houses. Table 1 gives results from an energy audit performed by the Pecan Street Research Institute and includes information on general housing characteristics. The statistics (average and standard deviation) refer to the neighborhood containing the 41 evaluated houses.

Table 1. Mueller houses' basic characteristics. (Rhodes et al. 2014)

Audit Field	Average	Median	St. Dev.
Year Built	2008	2008	0.7
Number of Levels	1.7	2	0.5
Conditioned Area (m^2)	192.5	192.1	50.0
A/C Capacity (kW)	10.6	10.6	2.8
A/C Efficiency (EER)	10.6	11	1.4
A/C Age	2008	2008	0.7
HVAC Duct R-value	6.8	6	1
Duct Leakage (%)	15.5	15	3.8

Model Form

The reduced-order model is an autoregressive with exogenous inputs (ARX) model, which means that it uses information from previous time steps to make estimates for the next time step. The formulation for the ROM derived by Cole *et al.* is given by

$$\begin{aligned} y_{i,j} &= a y_{i,j-1} + \sum_{k=0}^{2}\left[b_k DBT_{j-k} + c_k T_{i,j-k} \right] \\ &+ d\left[DBT_j \right]^2 + fDBT_j \cdot T_{i,j} + h_j \end{aligned} \tag{1}$$

where $y_{i,j}$ is the A/C electricity consumption for house i at time j, DBT_j is the outdoor dry bulb temperature at time j, $T_{i,j}$ is the thermostat set point for house i at time j, and parameters a-h are house-specific model coefficients (Cole, Powell, et al. 2013). The parameter h_j allows for a different constant term for each time period of the day, which accounts for other disturbances such as occupancy and solar irradiation. It was determined in the previous study that the ROM is accurate when it includes the last two time step values of the dry bulb temperature and thermostat set point. Accuracy was only marginally increased when the number of time steps was increased. Similar findings were found when the ROM was applied to this data set. This model was chosen because the inputs to the model (the house thermostat set points) and the outputs (the hourly air conditioning electricity consumption) make it useful for developing an A/C house energy management system through thermostat control.

Assumptions/Limitations

It should be noted that the model given in equation (1) is only valid during cooling periods. Also, because this is a data-driven model, care must be taken to not over-extrapolate beyond the temperatures and set points observed in the training dataset. Because of the form of the ROM, the physical characteristics of the houses do not need to be specified. However, the inputs into the model give an inherent assumption that A/C energy use is primarily driven by the thermostat set point and outdoor air temperature. In practice, the accuracy of this ROM had an adjusted R^2 value of 0.998 when compared with a detailed EnergyPlus model (Cole, Powell, et al. 2013).

Results

Reduced-Order Model Parameter Estimation Results

For each house, the disaggregated A/C energy and temperature data for the first 75% (38 days) of the data were used to regress the model coefficients in the ROM. Then, the remainder of the data was used to evaluate the accuracy of the derived model coefficients. A sample plot of the estimated energy values for the validation period is given for a sample house in Figure 2 on an hourly basis. For both the evaluated and training data, the estimated values closely align with the actual values. However, sharp increases in A/C use are dampened in the estimated energy values, and such changes in energy consumption are not always captured by the model. In addition, occasionally in the measured energy values there were zero energy hours during which the A/C was visibly off. Such events are not captured by the model estimations. There may be other unaccounted for events that resulted in a lower coefficient of determination. Although the R^2 value for this home was 0.752, the estimated values still strongly align with the measured values. There were two houses included in the data set which had stationary thermostat set points. These two houses were removed from this analysis because it is impossible to estimate the c and f coefficients in equation (1) if the thermostat set point remains constant. The mean R^2 value house for the remaining 41 houses during the evaluated period was 0.769 with a standard deviation of 0.080. The min and max R^2 values for the houses were 0.636 and 0.991. The mean difference between estimated and actual A/C electricity loads was 0.052 ± 0.026 kW hourly.

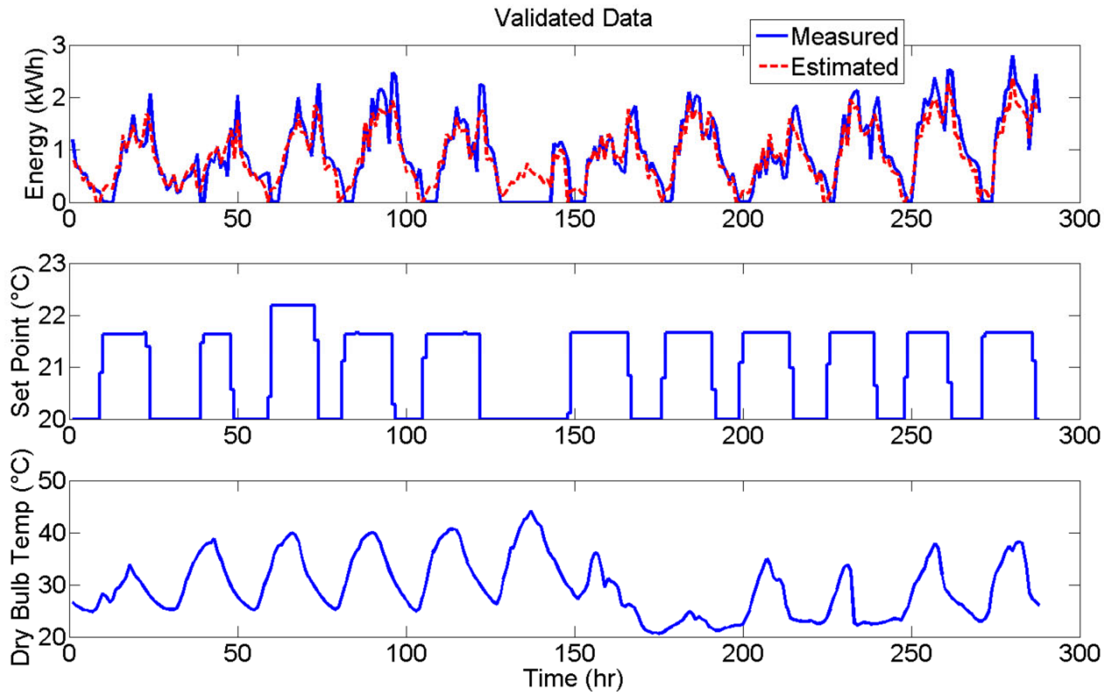

Figure 2. The top figure shows the hourly measured and estimated energy consumption from the A/C system for a single house. The middle plot shows the thermostat temperature set point, and the bottom plot shows the outdoor dry bulb temperature. The R^2 value for the top plot is 0.752, which is near the average R^2 for the 41 homes.

While there may be frequent deviations from actual values on an hourly basis, daily overall energy strongly aligns with the estimated values. Figure 3 shows the true and estimated total daily A/C energy consumption for the same house as above using the validation data set. The plot shows a high level of accuracy at the total daily level as seen in the comparable values. The reason for this may be that while the hourly estimates may not fluctuate as strongly as the actual measured A/C energy use, the general trends in correlation with the set points and outdoor air temperature are reasonably captured. The mean R^2 value for all the houses during the evaluated period was 0.967 for daily basis with a standard deviation of 0.044. The mean difference between estimated and actual loads was 0.272 ± 0.201 kWh for a day.

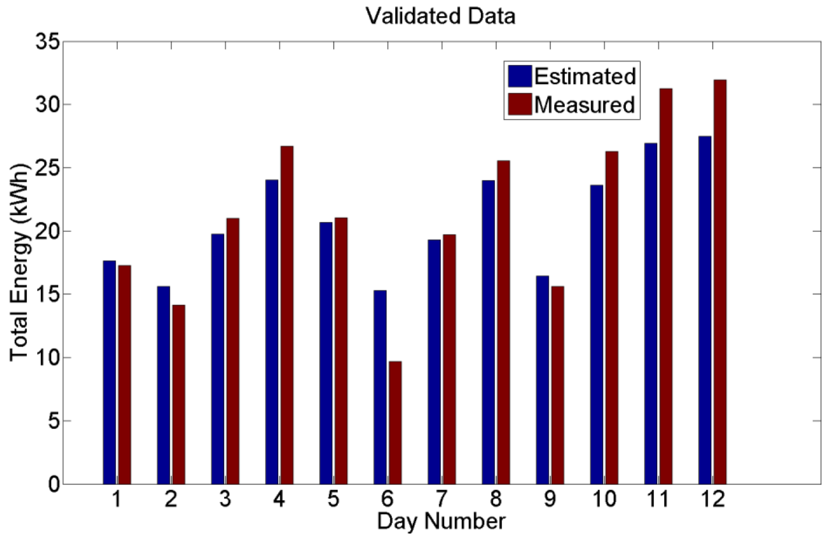

Figure 3. Daily total estimated vs. true energy for validation time periods for one house.

Model Predictive Controller

The ROMs were then used in a model predictive control (MPC) framework to determine the optimal thermostat set points that minimize peak electricity demand of the entire community through centralized control over the prediction horizon M of 11 hours (Cole, Rhodes, et al. 2013). MPC solves a limited-horizon optimal control problem that minimizes some objective function subject to a set of constraints. Only the first solution is implemented, then the time is advanced by one time step and the problem is re-solved. The formulation for the MPC to minimize peak energy is given by

$$J = \min_{T} z \tag{2}$$

Subject to

$$z \geq \sum_{i=1}^{N} y_{i,j} \qquad \forall j \tag{3}$$

$$y_{i,j} = f\left(y_{i,j-1}, T_{i,j-2:j}, DBT_{j-2:j}, j\right) \qquad \forall i, j \tag{4}$$

$$0 \leq y_{i,j} \leq maxLoad_i \qquad \forall i, j \tag{5}$$

$$lb_{i,j} \leq T_{i,j} \leq ub_{i,j} \qquad \forall i, j \tag{6}$$

where J is the objective function value, z is the maximum energy use of the community of houses over the time horizon of $j = t$ to $j = t + M$, i is the index for the houses, j is the index for the time, N is the total number of houses, $y_{i,j}$ is the A/C energy usage for house i at time j, f is the linear reduced-order model given by equation (1), $maxLoad_i$ is the maximum electricity consumption of the air conditioning unit of house i for one time step, and $lb_{i,j}$ and $ub_{i,j}$ are the lower and upper bounds, respectively, of the thermostat set point for house i at time j (Cole, Rhodes, et al. 2013).

In this control scheme, the thermostat set point (T) is the manipulated variable and the air conditioning electricity consumption (y) is the controlled variable. The outdoor dry bulb temperature (DBT) is the disturbance variable. Weather predictions for the dry bulb temperature (T) are assumed to be perfect in this paper. However, it has been shown that predictive controllers for building HVAC systems using a simple weather model can get within 1-2% of an optimal, perfect-prediction solution (Henze et al. 2004). It is assumed that occupants cannot override the set points chosen by the controller. Solving equations **Error! Reference source not found.-Error! Reference source not found.** leads to optimal thermostat set points ($T_{i,j}$) over the 12 hour time period for each house i. At 12 hours, the optimization problem can "see" the future peak and make early decisions to lower the peak through control actions. Then the time increases by one step and the next solution is found. This happens sequentially until the end of the evaluated time period of one day.

The upper and lower bounds of each house were found by using the thermostat set point data. The lower bound for each individual house was found by identifying the minimum set point throughout the entire data set for that house. This value was assumed to be the lower bound for the users in terms of thermal comfort. The upper bound was given two values depending on the time of day. During typical workday hours (8:00-17:00) the houses were assumed to be unoccupied. From the data set itself, it was unclear how to reveal which homes were occupied during work hours so the same assumption was applied to all homes. Accuracy could be increased if a house survey or an appropriate sensor reported the typical occupancy hours. The maximum set point throughout the entire data set for each individual house was set to be the maximum set point during unoccupied hours. During occupied hours, the maximum set point was assumed to be the minimum set point plus 2.22 °C (4 °F). For houses that had a narrower minimum and maximum set point band than 2.22 °C (4 °F), the maximum set point value was used for all hours during the day. In this way, the range of allowable set points for each house never extrapolated beyond the region where the ROM was trained. An example of the upper and lower thermostat set point bounds for a single house is given in Figure 4.

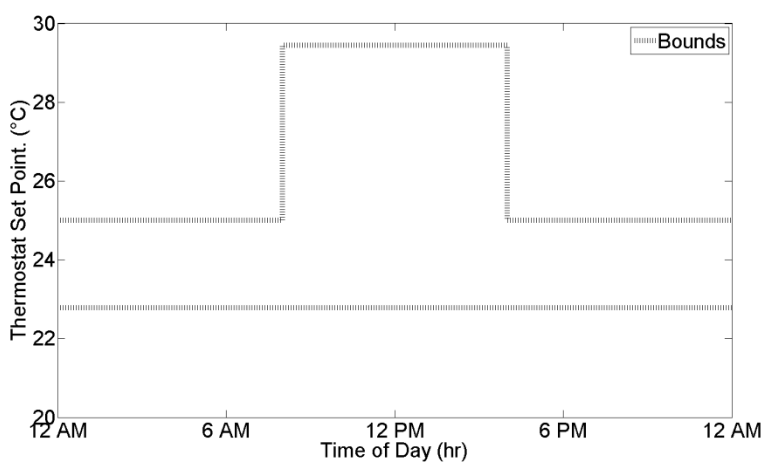

Figure 4. The upper and lower bounds for one house. The house is unoccupied from 8:00-17:00. The desired set point for this house would optimally be at 22.7°C (73°F). These bounds constitute lb and ub in (6).

Lastly, it was noticed in the previous study that occasionally houses have undersized air conditioning units or thin building envelopes. Therefore, at low thermostat set points and high outdoor temperatures, the air conditioner could not meet the load required to keep the indoor temperature below the upper bound without violating (6). Therefore, the upper bound for (6) was given as a soft constraint with a high penalty for violation.

The 41 houses were used in this centralized controller. Recall that the controller acted to minimize the total peak value of the neighborhood for each day. An example of the controller actions taken to minimize the peak consumption for one house is shown in Figure 5. For this house, the thermostat set point is lowered during typically unoccupied hours prior to the late afternoon when temperature is highest. The lower set point helps to pre-cool the house and store the thermal energy so that during peak hours the set point can be held at the upper bound.

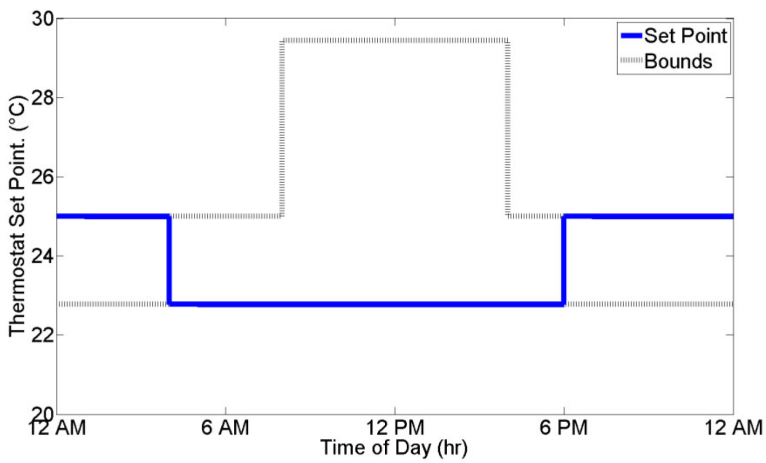

Figure 5. Sample optimal control of one house for one day. The change in the set point temperature pre-cools the house to stagger when the A/C unit turns on so that the peak energy for the neighborhood is minimized.

The overall benefit of utilizing the inherent thermal mass of each building is seen in Figure 6, which displays the energy consumed by the A/C for all 41 houses during one day. The load is substantially leveled during peak A/C consumption times. The centralized controller is able to take advantage of passive thermal energy storage by pre-cooling the thermal mass of the houses. Therefore houses that have thinner thermal envelopes or tight thermostat constraints can still run during peak times while those with large thermal envelopes can slowly heat up. In both cases the thermal comfort bounds of each house are not violated.

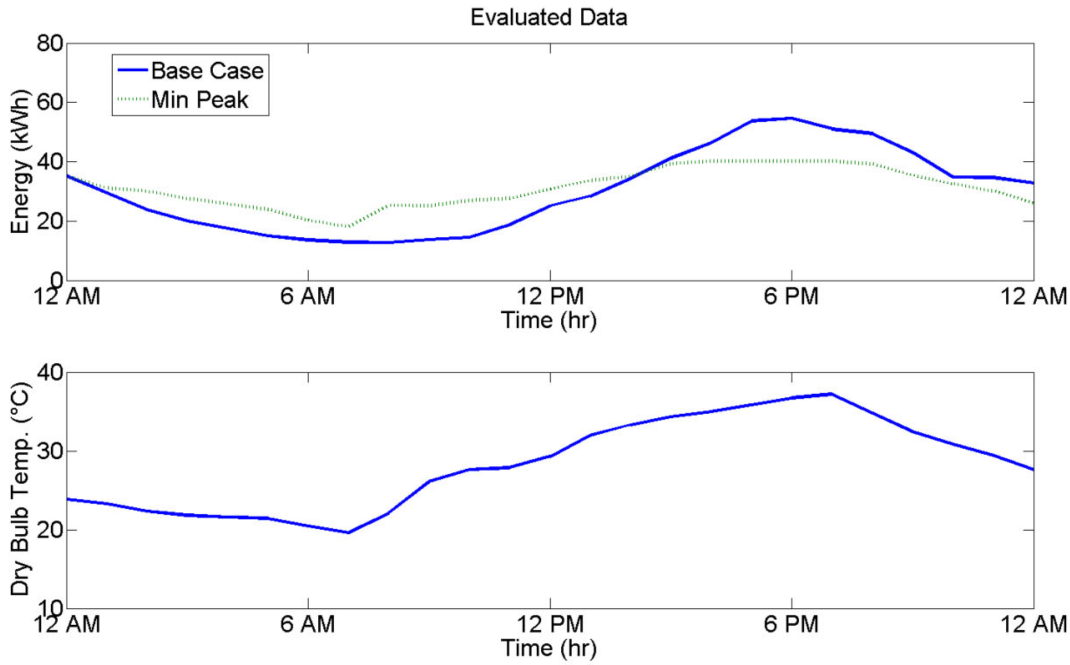

Figure 6. The upper plot shows the total A/C energy consumption of the 41 houses for June 2. The lower plot shows the outdoor dry bulb temperature for the day. The "Base Case" refers to how the sum of the A/C units actually operated. The "Min Peak" refers to the centralized control strategy implemented to level the peak load.

There were 11 consecutive days that all houses had complete data that could be used for comparison. Table 2 shows the comparison of the measured energy consumption (the base case) compared with the simulated minimum peak control scheme (the minimum peak case). In general, the minimum peak simulation consumes less energy than the base case. The average set-point temperature for the base case is about 0.4°C (~1°F) higher than the set points in the controller case, which indicates that some of the houses are on average cooler than in the base case. It is expected that the minimum peak control scheme should use more energy to pre-cool the house because some of that energy is lost to the environment while the house warms up. However, the total energy values are closer together because the minimum peak controller works at the thermal comfort bounds estimated for the consumers and thus saves energy by operating near or at the boundaries. If the bounds were tightened during occupied hours the results might change. Another possible reason there is lower energy consumption when the average set point temp is lower is because the A/C efficiency is a function of outdoor dry bulb temperature. Thus shifting more of the cooling load to the morning may improve the efficiency. In this comparison the trade-off between total A/C energy use and peak power reduction is not strong, but in general the total energy use should increase in exchange for the ability to lower the peak load.

The greatest benefit of the controller was in reducing the peak load. In all cases the peak load was significantly reduced by the controller action, with an average reduction of 27% compared to the base case. Table 2 displays the overall results of the control optimization. The centralized controller leverages the thermal mass of each house, which resulted in a significant cut to the peak A/C power. Through balancing the total load of the houses using pre-cooling, the controller staggers when units turn on. The central controller is thus able to shift the peaks of each individual house and successfully level the overall load. This operational scheme was not explicitly stated in the controller formulation. Rather, it is the result of the optimization and one of the benefits of centralized control.

Table 2. Total A/C energy, peak power, and average set point temperature for the base case and the peak minimization case ("Min Peak") using centralized control. The "Max Temp" column shows the maximum daily outdoor temperature. On average, the A/C peak power was reduced by 27%

		Total A/C Energy (kWh)			A/C Peak Power (kW)			Average Set Point Temp (°C)	
Day	Max Temp	Base Case	Min Peak	Savings	Base Case	Min Peak	Savings	Base Case	Min Peak
1	37.02	768.81	689.91	10.3%	58.38	39.98	31.5%	25.02	24.67
2	33.00	721.13	744.34	-3.2%	54.55	40.21	26.3%	25.02	24.74
3	37.20	685.33	692.10	-1.0%	54.62	39.77	27.2%	25.02	24.72
4	35.72	691.77	689.61	0.3%	50.98	39.67	22.2%	25.02	24.72
5	35.67	690.11	683.30	1.0%	54.37	38.90	28.5%	25.00	24.73
6	39.12	746.48	687.18	7.9%	58.55	39.89	31.9%	24.91	24.72
7	33.28	683.97	730.64	-6.8%	54.22	40.34	25.6%	24.92	24.66
8	35.05	685.21	702.33	-2.5%	53.33	40.23	24.6%	25.19	24.69
9	34.27	725.39	685.02	5.6%	58.55	39.36	32.8%	25.34	24.73
10	37.27	729.98	702.97	3.7%	56.09	40.31	28.1%	25.36	24.72
11	35.03	702.43	697.64	0.7%	53.47	39.82	25.5%	25.35	24.71

Conclusions and Future Work

In this work a ROM that was developed for EnergyPlus models was applied to 41 real houses using smart meter and smart thermostat data. The ROM form was able to accurately predict energy usage on the real houses. The ROM predicted total daily energy with a mean R^2 value of 0.967 for all 41 houses with a mean difference between estimated and actual loads of 0.272 ± 0.201 kWh. On the hourly level, the ROM predicted energy consumption with an R^2 of 0.724. The mean difference between estimated and actual loads was 0.052 ± 0.026 kWh on an hourly basis. The availability of data limited the study to just two months.

It has been verified that through an aggregation of houses, load forecasting and peak energy reduction techniques can be evaluated at the community level. A community-level model predictive control (MPC) scheme was implemented using the ROMs of the 41 homes (2 homes were removed due to insufficient data). In the control demonstration, it was shown that a centralized MPC controller could be used to reduce the peak load by pre-cooling homes and staggering the time A/C units turned on. On average the MPC was able to reduce the peak load by 27% (15 kW) for the group of 41 houses without significantly increasing electricity consumption.

Future work includes increasing the accuracy of the ROM. An expansion of the available data set will create the opportunity for more detailed investigation of the robustness of the ROM. Furthermore, while a ROM removes a significant amount of physical insights by lumping into regressed coefficients, there might be a limited physical interpretation of some of the parameters. Because the ROM, which was developed from an EnergyPlus model, was successfully implemented on an actual home, it may be possible to perform a parametric analysis to attach physical meaning to the ROM coefficients. This will make it possible to evaluate how changes in the house, such as increased insulation, can change the A/C usage and possibly suggest retrofits for current houses.

Acknowledgements

This work was supported by the Pecan Street Research Institute (a 501(c)3 nonprofit public-private partnership in Austin, Texas), the United States Department of Energy, and the University of Texas at Austin. This material is also based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1110007. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Anthony, Kelman. 2011. “Bilinear Model Predictive Control of a HVAC System Using Sequential Quadratic Programming.” In , edited by Bittanti Sergio, 9869–74. doi:10.3182/20110828-6-IT-1002.03811. http://www.ifac-papersonline.net/Detailed/50719.html.

Christensen, Craig, Dave Roberts, Ben Polly, Marcus Bianchi, Neal Kruis, and Ren Anderson. 2010. “Building Energy Simulation Accuracy” July 21, National Renewable Energy Laboratory.

Cole, Wesley J., Kody M. Powell, Elaine T. Hale, and Thomas F. Edgar. 2013. “Reduced-order Residential Home Modeling for Model Predictive Control.” *Energy and Buildings*, June.

Cole, Wesley J., Joshua D. Rhodes, William Gorman, Krystian X Perez, Michael E. Webber, and Thomas F. Edgar. 2013. “Community-scale Residential Air Conditioning Control for Effective Grid Management.” *Applied Energy*, October.

Gouda, M. M., S. Danaher, and C. P. Underwood. 2002. “Building Thermal Model Reduction Using Nonlinear Constrained Optimization.” *Building and Environment* 37 (12): 1255–65. doi:10.1016/S0360-1323(01)00121-4.

Henze, Gregor P., Doreen E. Kalz, Clemens Felsmann, and Gottfried Knabe. 2004. “Impact of Forecasting Accuracy on Predictive Optimal Control of Active and Passive Building Thermal Storage Inventory.” *HVAC&R Research* 10 (2): 153–78. doi:10.1080/10789669.2004.10391097.

Karmacharya, S., G. Putrus, C. Underwood, and K. Mahkamov. 2012. “Thermal Modelling of the Building and Its HVAC System Using Matlab/Simulink.” In *2012 2nd International Symposium on Environment Friendly Energies and Applications (EFEA)*, 202–6. doi:10.1109/EFEA.2012.6294078.

Lomas, K.J., H. Eppel, C.J. Martin, and D.P. Bloomfield. 1997. “Empirical Validation of Building Energy Simulation Programs.” *Energy and Buildings* 26 (3): 253–75. doi:10.1016/S0378-7788(97)00007-8.

Malisani, Paul, Francois Chaplais, N. Petit, and D. Feldmann. 2010. "Thermal Building Model Identification Using Time-scaled Identification Methods." In *2010 49th IEEE Conference on Decision and Control (CDC)*, 308–15. doi:10.1109/CDC.2010.5717975.

Perez, Krystian X, Wesley J. Cole, Joshua D. Rhodes, Abigail Ondeck, Michael Baldea, Michael E. Webber, and Thomas F. Edgar. 2014. "Nonintrusive Disaggregation of Residential Air-Conditioning Loads from Smart Meter Data." *Energy and Buildings*, Under Review. February 2014.

Rhodes, Joshua D., Charles R. Upshaw, Chioke B. Harris, Colin M. Meehan, David A. Walling, Paul A. Navrátil, Ariane L. Beck, et al. 2014. "Experimental and Data Collection Methods for a Large-scale Smart Grid Deployment: Methods and First Results." *Energy* 65 (February): 462–71. doi:10.1016/j.energy.2013.11.004.

U.S. Department of Energy. 2014. "EnergyPlus Energy Simulation Software." http://apps1.eere.energy.gov/buildings/energyplus/.

———. "eQuest Website." http://www.doe2.com/equest/.

Wattles, Paul. 2011. *ERCOT Demand Response Overview & Status Report*. http://www.ercot.com/content/meetings/dswg/keydocs/2011/0830/3_ERCOT_presentation_workshop_083011.pdf.