ABSTRACT

Title of Document:	**PERFORMANCE AND APPLICATIONS OF RESIDENTIAL BUILDING ENERGY GREY-BOX MODELS**
	Michael James Siemann, Doctor of Philosophy, 2013
Directed By:	Professor Jungho Kim, Department of Mechanical Engineering

The electricity market is in need of a method to accurately predict how much peak load is removable by directly controlling residential thermostats. Utilities have been experimenting with residential demand response programs for the last decade, but inconsistent forecasting is preventing them from becoming a dependent electricity grid management tool. This dissertation documents the use of building energy models to forecast both general residential energy consumption and removable air conditioning loads.

In the models, complex buildings are represented as simple grey-box systems where the sensible energy of the entire indoor environment is balanced with the flow of energy through the envelope. When internet-connected thermostat and local weather data are inputs, twelve coefficients representing building parameters are used to non-dimensionalize the heat transfer equations governing this system. The model's

performance was tested using 559 thermostats from 83 zip codes nationwide during both heating and cooling seasons. For this set, the average RMS error between the modeled and measured indoor air temperature was 0.44°C and the average daily ON time prediction was 1.9% higher than the data. When combined with smart power meter data from 250 homes in Houston, TX in the summer of 2012 these models outperformed the best traditional methods by 3.4 and 28.2% predicting daily and hourly energy consumption with RMS errors of 86 and 163 MWh. The second model that was developed used only smart meter and local weather data to predict loads. It operated by correlating an effective heat transfer metric to past energy data, and even further improvement forecasting loads were observed.

During a demand response trial with Earth Networks and CenterPoint Energy in the summer of 2012, 206 internet-connected thermostats were controlled to reduce peak loads by an average of 1.13 kW. The thermostat building energy models averaged forecasting the load in the 2 hours before, during, and after these demand response tests to within 5.9%. These building energy models were also applied to generate thermostat setpoint schedules that improved the energy efficiency of homes, disaggregate loads for home efficiency scorecards and remote energy audits, and as simulation tools to test schedule changes and hardware upgrades.

PERFORMANCE AND APPLICATIONS OF RESIDENTIAL BUILDING
ENERGY GREY-BOX MODELS

By

Michael Siemann

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:
Professor Jungho Kim, Chair
Professor Steven Gabriel
Research Associate Professor Yunho Hwang
Professor Michael Ohadi
Professor Reinhard Radermacher

Acknowledgements

more. The countless discussions and arguments in which we engaged in were invaluable to the progress and quality of this work.

I would like to thank many of my friends and engineering classmates for their support over the years. In particular, I would like to thank Adam Eshleman, Mary Larson, and Mark Patrick for helping me determine if ideas presented in this dissertation were coherent. I am also especially grateful to have loving parents, Art and Mary, who sparked my interest in science at a young age and taught me the value of hard work and being a good person.

Table of Contents

List of Tables

List of Figures

Note that many figures in this dissertation use a line to represent non-continuous data. Plotting the specific data points would be more correct, but was more difficult to read.

x

Nomenclature

Acronyms

AC	=	Air conditioning
ASHRAE	=	American Society of Heating, Refrigerating and Air-Conditioning
CDD	=	Cooling Degree Day
DR	=	Demand Response
ERCOT	=	Electric Reliability Council of Texas
GA	=	Genetic Algorithm
HVAC	=	Heating, Ventilation, and Air conditioning
ISO	=	Independent System Operator
RMS	=	Root Mean Square
TMY	=	Typical meteorological year

Symbols

A	=	area
α	=	thermal diffusivity
β	=	buoyancy coefficient of expansion
C	=	thermal capacitance
c	=	dimensionless coefficient
cp	=	specific heat
Δx	=	wall grid spacing
Δt	=	time step
fr	=	fenestration ratio
g	=	gravitational constant
$HVAC_{mode}$	=	HVAC Status (ON/OFF)
h	=	convective heat transfer coefficient
k	=	thermal conductivity
l	=	length
$\dot{m}$	=	mass flow rate
μ	=	dynamic viscosity
ν	=	kinematic viscosity

Ø	=	heat transfer
$\dot{Q}$	=	heat transfer
R	=	thermal resistance
ρ	=	density
T	=	temperature
t	=	time
u	=	wind speed
V	=	volume

Subscripts and Superscripts

a	=	indoor air
cond	=	conduction
i	=	grid point index
m	=	interior thermal mass
n	=	last wall grid point
∞	=	outdoor

Several terms are used thought this dissertation that may only be defined in one section. To increase the independence of each chapter, definitions of common terms are presented in the Glossary on page 152.

Chapter 1: Introduction and Contributions

1.1 Introduction

This dissertation documents the derivation and performance of two building energy models that can change the way residential loads are predicted by the energy market. These grey-box models utilize data from new technology in the residential sector, specifically internet connected thermostats, hyper-local weather sensor networks, and smart power meters, instead of on-site measurements. The building heat transfer is modeled in a way to yield quick computation and minimal data storage so it can scale to millions of homes for use by utilities and retail electricity independent system operators (ISO). Applications that utilize the building models are also presented in this dissertation, including thermostat control demand response forecasting, an energy efficiency simulation engine and control techniques, and remote energy auditing. This chapter presents the intellectual contributions made to the scientific community and the anticipated benefits for the energy industry and homeowners. Chapter 2 reviews the literature and current practices in building energy modeling, residential demand response, and load forecasting, as well as the motivation for this work. Chapter 3 presents the derivation, evaluation, and performance of the two building energy models predicting general loads. Chapter 4 explores applications that have been developed and tested using these building energy models, and Chapter 5 presents the conclusions and directions for future work. Also note a glossary of terms used throughout the work is found at the end of the

dissertation.

1.2 Intellectual Contributions

Electricity utilities and other members of the energy market need a way to accurately predict how much removable peak load is available by controlling residential thermostats. Studies found in the literature, results presented by utilities, and even tests performed as part of this study all reported inconsistent removed capacity because of the multitude of variables present. A building energy model was developed with new technology to improve this. Another building energy model was developed to forecast loads when internet-connected thermostat data is not available. Both these models are also used as tools to provide analytics to individual customers and members of the energy market.

1.2.1 Thermostat Building Energy Model

Speed

A grey-box building energy model was developed using internet-connected thermostat, local weather station, and smart power meter data. The model was designed to accurately forecast individual and aggregated loads quickly. It was observed that the building model aspects of daily load forecasting were executed on the order of milliseconds using a standard desktop computer. Additionally with the performance of cutting-edge cloud databasing, an analysis of the situation could be performed and customized commands could be generated and sent to hundreds of internet-connected thermostats to initiate a demand response event in several minutes.

Data

The building energy model was designed to require a minimal amount of data to be stored. The thermodynamic characteristics of a house are represented by 15 double-precision numbers. Add in station identification and location values and the storage required for a model is less than 200 bytes. Without the equations and normalizing factors this data is also meaningless, so security is less of a concern.

Range

Test were performed documenting the performance of the model using data from 559 thermostats during heating (January 2013), cooling (July 2012), and a shoulder season (May-June 2012). The houses tested were from 83 different zip codes and five of the seven ASHRAE climate zones. The only statistically significant factor observed that influenced a model's ability to forecast was the thermostat setpoint schedule, and this is easily adjusted. The range in locations, climate, and HVAC operation successfully modeled in this study suggests that this model can be applied to any house nationwide.

Accuracy

In the 559 thermostat set the model averaged monthly ON time predictions only 1.9% higher than the measured data and an indoor temperature RMS error of 0.44°C. When combined with smart meter data, the model predicted the daily RMS total energy of 250 additional houses in Houston, TX from July 1 to September 30, 2012 3.4% better than the highest performing traditional prediction technique at an RMS error of 86 MWh. The hourly RMS error of 163 MWh was 28.2% better as well. This performance was based off of actual weather data, but when weather forecasts are used the model improved the daily prediction by 12.7% and hourly prediction by

5.4% with RMS errors of 89 and 191 MWh. The model averaged forecasting the load of nearly 200 thermostats to within 5.9% during demand response events. These encompassed the 2 hours before, during, and after three demand response tests in the late summer of 2012.

1.2.2 Smart Meter Only Model

The smart meter only building energy model was developed using only smart meter and local weather station data. The model may be unable to predict removable load from the control of thermostats, but it was the most accurate model tested at predicting general loads. Effective heat transfer from the outdoor to the indoor environment serves as metric to drive this model.

Speed

Though not tested, it can be assumed this model will compute faster than the thermostat building energy model because it is governed by only one overall heat transfer equation (Equation 3.52).

Data

The same performance involving the data required for the thermostat building energy model is repeated for this model

Range

Only 250 houses were tested in Houston, Texas. While not as large and diverse as the set tested by the thermostat building energy model, a wide range of houses and conditions was still covered.

Accuracy

This was the highest performing model tested at predicting the general daily and

hourly load throughout the summer of 2012 in Houston, TX. This model improved on the best tested traditional technique by 6.7% for the daily prediction and 39.6% for the hourly prediction with RMS errors of 83 and 137 MWh. It also improved the hourly predictions by 8.4% when weather forecasts were used with an RMS error of 185 MWh.

1.2.3 Tools

Several novel techniques and tools were developed in the process of creating these models. These can be applied to other energy models and applications not presented.

<u>Disaggregating Load</u>

Two methods were developed to disaggregate HVAC load from the total house electricity load. The first method used internet-connected thermostat and local weather station outdoor air temperature data to approximate a power curve. From the power curve the HVAC load during any duration of when it was ON was estimated. The second method correlated energy consumption to the heat transfer metric. This metric was used as a filter to identify the power signatures of base and HVAC load in smart meter data, and those signatures were used to further disaggregated the load. Simulations with the thermostat building energy model were performed to isolate how much energy was consumed by each major energy flow to maintain indoor comfort. The weather data was adjusted to separate the energy consumption attributed to solar loading, infiltration, and general heat diffusion through the envelope.

<u>Thermostat Schedules</u>

Three mechanisms were developed to reduce the energy consumed by the HVAC

system maintaining indoor comfort. Smart setback allows homeowners to set a desired indoor air temperature schedule, not simply a setpoint schedule. The developed algorithm determines new setpoints that will meet the desired temperature most efficiently. Energy savings is realized by instilling more aggressive setbacks in schedules that previously had no or moderate setbacks. Setpoint smoothing starts setback periods earlier if the building model forecasts the indoor air temperature will not increase 1°C above the current setpoint. This mechanism reduces the time spent controlling to indoor temperatures further from the thermal equilibrium with the outdoor environment. Precooling takes advantage of the higher cooling efficiency of colder outdoor air temperatures. These mechanisms are explained using cooling examples but the same logic can be applied for heating. Also variable rate price schedules can be added to reduce the overall cost of maintaining indoor comfort.

1.3 Anticipated Benefits

1.3.1 Energy Markets

The thermostat building energy model is a tool that can be used to evaluate the available capacity of controlling thermostats for peak load reduction. Simulations can be performed with these models to determine the control strategies that best meets the needs of the utilities, retail electricity providers, energy traders, aggregators, ISO, etc. and also the individual homeowners. The appropriate balance between capacity and customer comfort can be determined before the systems are relied upon when managing the grid. With accurate weather forecasts the thermostat and smart meter only models can be used to forecast residential loads for the upcoming days. This can reduce the risk members of the energy market need to take in fulfilling their

commitments.

1.3.2 Homeowners

The building energy models give homeowners tools to reduce their energy consumption and bill. The mechanisms developed to reduce energy consumption by manipulating the thermostat setpoints can save significant energy and money, especially if they convince a user to be more aggressive with setbacks. These can be used to avoid or reduce consumption during expensive periods of variable rate plans. The models can be used to provide a Home Efficiency Scorecard where customers can learn about their energy consumption. A breakdown of different loads and rankings can help identify what issues should be targeted first. The models can be used to forecast utility bills and offer strategies to control a home given a budget. The payback period of retrofit upgrades can also be tested to help apprehensive homeowners look beyond the initial costs. If the members of the energy market use the tools effectively homeowners can see reduced rates and a more reliable electricity grid.

Chapter 2: Background

2.1 Building Energy Modeling

Buildings are complex systems that are influenced by numerous inherent and controllable factors. Consequently, researchers have developed many modeling techniques to predict building behavior (e.g, Zhao and Magoulè 2012). With regard to predicting building energy consumption, the methods used to model the building can be classified as white-box, black-box, or grey-box systems. White-box systems are based on detailed knowledge of building's physical characteristics. They require a complete understanding of the system, and often involve extensive fine-tuning, experience, and computational effort to be accurate. Examples of building white-box modeling software are TRNSYS, DOE2, and EnergyPlus. Extensive work has been done to highlight their advantages and identify their shortcomings (Lowry and Lee 2004). Black-box models require minimal understanding of the physical system. Whether they are statistics-based or employ complex mathematical algorithms, these empirical models require extensive historical training data and perform poorly when the input data are outside the bounds of the training set. The most popular black-box models in energy are Artificial Neural Networks (ANN). Li et al. (2009) successfully reproduced the energy consumption data generated from the white-box DeST model using four different ANN techniques. When compared to the data measured in a test office building in Brazil, Neto and Fiorelli (2008) were able to outperform the 13% energy prediction error observed using EnergyPlus modeling with their simple ANN technique that had only 10% error. With ANN and other black-box models, a major

requirement for successful energy predictions is diverse independent inputs in the training set, especially the weather conditions. Black-box model's perform poorly when having to extrapolate to inputs not seen in the training set. Henze et al. (2004) used several different collections of weather data to build simulation models with the goal of optimizing the energy consumption of a simulated office building. When compared to a simulation where the actual weather data was used to create a load profile, running a simulation with a model trained using the last 30 days outperformed models using only yesterday's data, the last week's data, and an hourly random walk temperature.

Grey-box models bridge the gap between white-box and black-box systems (see Oussar and Dreyfus 2001). Wu and Sun (2012) introduced building physics to the autoregressive moving average method and observed a 1% RMS error in the cooling demand of a 4-story California office building. A simpler modeling approach, employing the popular thermoelectricity analogy, was taken by Braun and Chaturvedi to model office buildings in Chicago, IL (2002). Figure 2.1 shows the circuit model of the five different types of structures considered: external walls, ceiling/roof, floor, internal walls, and windows. The Ts, Rs and Cs in the figure denote temperatures, resistances, and capacitances with the subscripts e, c, f, i, w, a, z, g representing external, ceiling, floor, internal, window, ambient air, zone air, and ground, respectively.
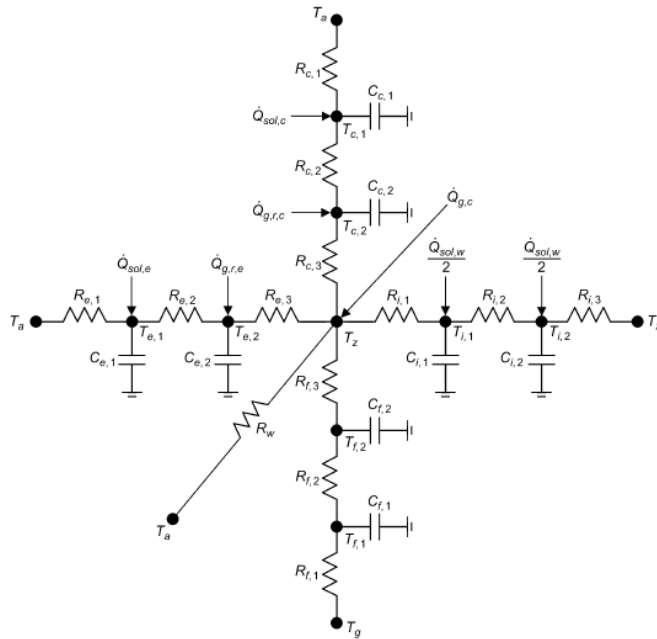
Figure 2.1: Thermal network for overall building model (Braun and Chaturvedi 2002).

Their inverse grey-box model was able to achieve 8.6% RMS error in the predicted sensible load using only 2 weeks of training data in the summer. Figure 2.2 shows how well the model was able to match the 14-day training data cooling load.
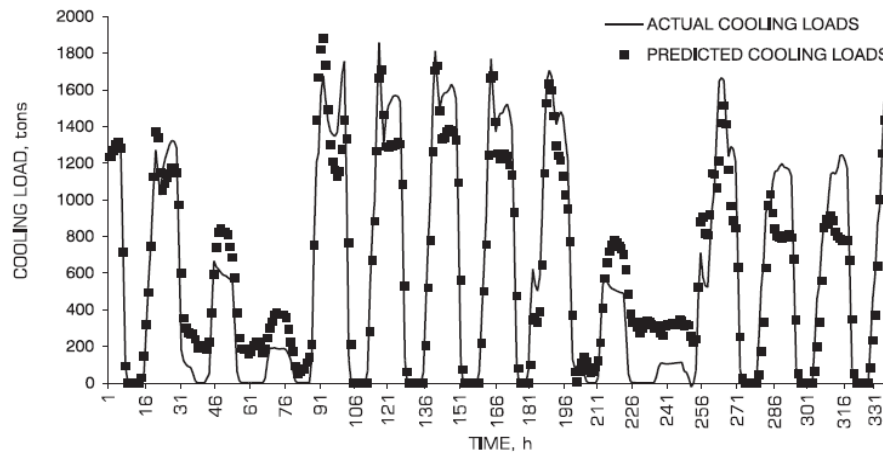
Figure 2.2: Comparison of actual and predicted sensible cooling loads for training period (Braun and Chaturvedi 2002).

10

Grey-box models appear to be the best balance between resources and effort in predicting energy consumption relative to white and black-box models, but are still only being used for commercial buildings. Residential buildings have yet to see the widespread adoption of any type of modeling for two main reasons: predictability and benefit. While commercial buildings can be difficult to model, houses tend to respond to external (weather) and internal (thermostat setpoint) changes even more radically. Due to their lower thermal mass, when compared to commercial buildings, along with a lack of distributed control, houses are too dynamic for black-box models. Irregularities in the occupant patterns and construction also make it difficult and cost prohibitive to generate accurate white-box models.

The weather data used to characterize the conditions on the outside of the building envelope is equally as important as the physics governing the model in simulations. The standard for most models is typical meteorological year (TMY) data which is assembled using weather station data from locations across the United States and its territories. TMY3 (2013), the latest iteration of data, gives hourly averages for 1020 locations for an entire year. The dataset includes dry bulb air temperature, relative humidity, wind speed, and solar load in several values (global horizontal irradiance is the most useful for energy simulations because it is comprised of both direct and diffuse irradiance). Figure 2.3 gives a map of the locations, but note that the majority of the stations (class II and III) are generated with lower quality measuring equipment and contain gaps in the period of record (Wilcox and Marion 2008).
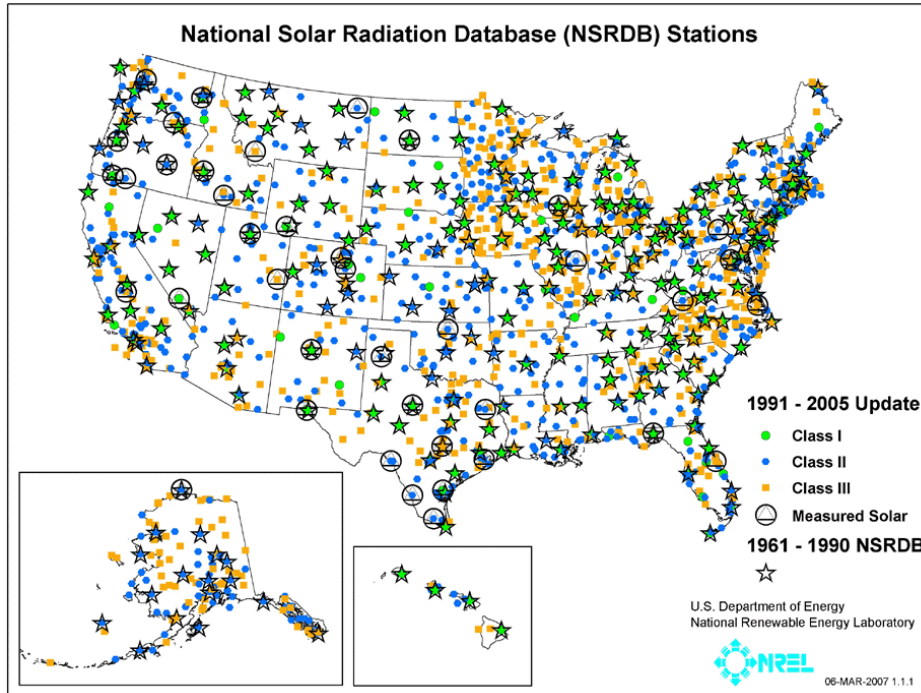
Figure 2.3: Map of TMY3 locations grouped by quality of station data (Wilcox and Marion 2008).

While TMY data is useful for predicting seasonal averages, it fails to provide a resource for testing weather conditions that are not typical. Heating and cooling systems are often sized for the extreme weather conditions a building might encounter, not the average seasonal conditions. Demand response events often occur when building HVAC use increases significantly from the norm and these cannot be forecasted using averages. As an example, Figure 2.4 shows how the electricity load of 250 homes in Houston, TX differed from the average load profile (derived from the entire summer). The solid black line plots the load every 15 minutes, the dashed grey line plots the average, and the shaded regions reflect how much they differ. Note that this figure, and many figures in this dissertation, uses a line to represent non-continuous data. Plotting the specific data points would be more correct, but was

more difficult to read.
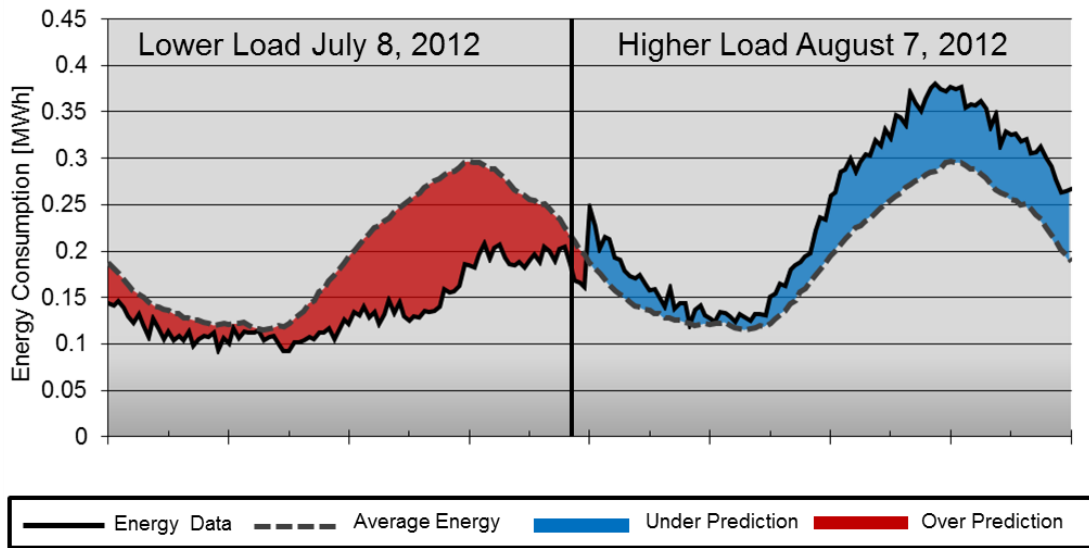
Figure 2.4: Comparing seasonal average load data to days with lower and higher energy consumption for an aggregate of 250 Homes in Houston, TX.

If TMY weather data is not available for a location, simulated data can serve as an alternative (Degelman 1991, 2004). The statistical-based weather set generator tested by Degelman matched the average conditions well but failed to predict extremes. The national weather service collects data from 1500 stations nationwide but most do not record solar irradiance. Also because most of these weather stations are at airports, where geographical conditions can vary significantly from the cities they service, a disconnect is observed between these conditions and what is actually occurring at the location of the building. An improvement on this is the Earth Networks WeatherBug (2013) weather station network. It is comprised of over 8000 stations nationwide, often in residential and commercial locations, and contains the measurements needed for building modeling (see Figure 2.5).

Figure 2.5: The Earth Networks WeatherBug weather station network in the USA (Earth Networks 2012).

2.2 Shifting Residential Peak Loads for Demand Response

Peak demand for electricity on hot summer afternoons is common nationwide primarily because of the dynamic loads of air conditioning. The air conditioning load on the grid increases in the afternoon when outdoor temperatures and solar load increase. This is in contrast to relatively constant daytime load profiles such as commercial lighting and industrial processes that make up the majority of the base load. Figure 2.6 shows the afternoon peak by plotting the residential load along with disaggregated air conditioning load. The difference between the two power signals also increases in the afternoon when other time dependent loads such as lighting and appliances are used more heavily.

14
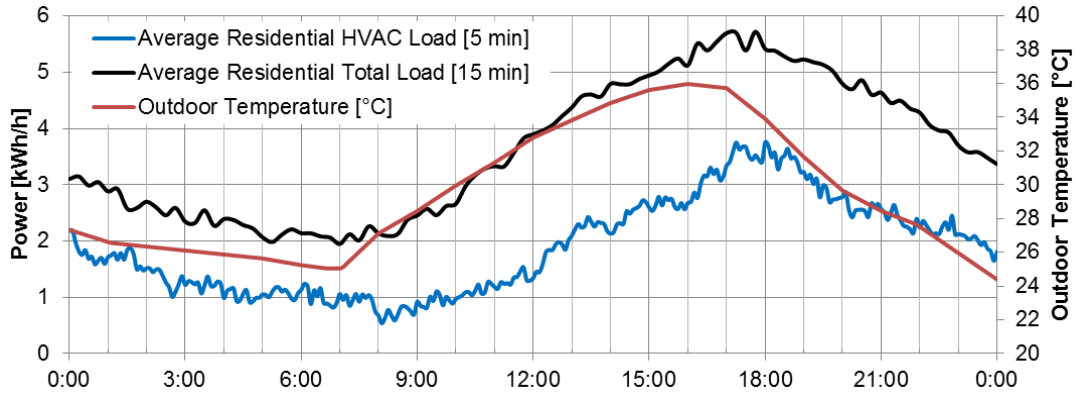
Figure 2.6: Average residential electrical power load in Houston, TX on August 24, 2012 from 100 homes. The HVAC load is calculated using thermostat runtime data and effective power curves.

To compensate for peak loads, utilities must import or generate additional capacity and/or reduce the demand. Failure to match supply and demand will result in localized brownouts or blackouts. Importing power is often at a high cost premium, as is switching in peak capacity generators. These generators, often referred to as *peakers*, spend the majority of their operational life OFF or on standby and yet still require nearly the same maintenance and cost to run as uninterruptible base load plants. Also combine that with the inability of many utilities to construct additional generation, transmission, and distribution fast enough to accommodate projected demand growth, to see how important reducing the demand will become. Strategies that reduce the demand for power during peak times by shifting or curtailing its use are referred to as *demand response*. Because the residential sector significantly contributes to the peak load, methods targeting residential customers have been tested to reduce it. The two major strategies involve offering incentives to shift or reduce their load or directly controlling devices and appliances.

15

Incentives to get customers to take an active role in reducing residential peak loads often come in the form of time-varying pricing. Newsham and Bowker (2011). outline the four main rate styles: time-of-use (TOU), critical peak pricing (CPP), real time pricing (RTP), and peak time rebates (PTR). TOU, CPP, and RTP all have higher rate prices during peak times or demand response events, and PTR pay customers for the power they did not use during peaks. In the summer of 2004, Herter and Wayland (2009) were able to test load reduction using a CPP prices of $0.50/kWh and $0.68/kWh on 483 houses in California. When compared to the baseline of non-event days, the averaged peak power reduction on 12 event days was 0.07 kW (5.1%) for both price groups. The average response of the 12 event days is plotted in Figure 2.7. Note though the consumption reached an average of 0.10 kW above normal consumption after the event period; this period will be referred to as the *recovery* phase of a demand response event for the duration of this dissertation.
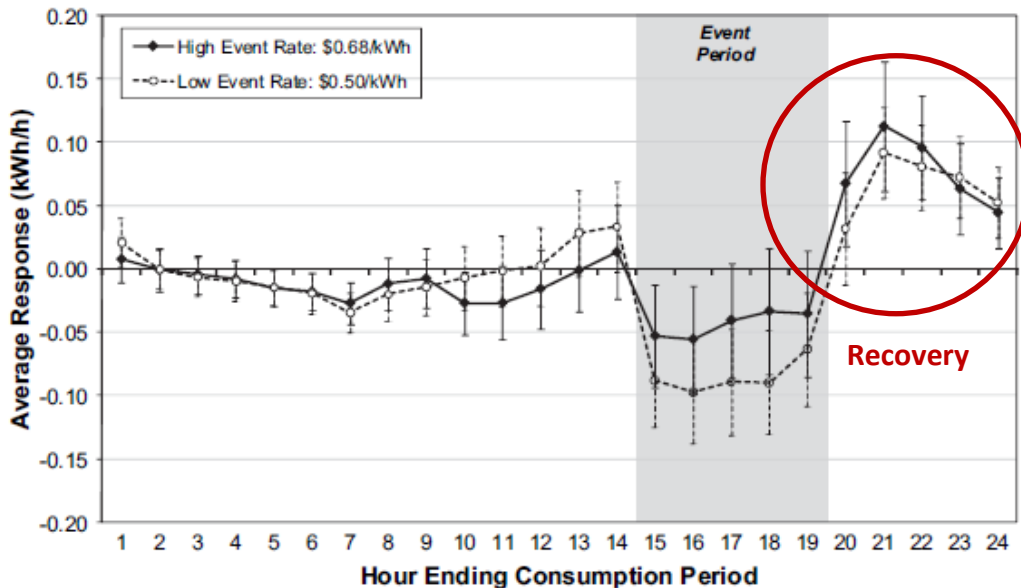
Figure 2.7: Residential demand response: high vs. low-ratio CPP tariff (modified from Herter 2009).

One drawback with all these price approaches is they require power meters that store or correlate the energy usage with the time of day; though, utilities and the US government are investing to increase the adoption and installation of these smart power meters (DOE 2012). Another barrier is that the customers have to take an active role in the process. Whether turning off appliances or adjusting the setpoints on a thermostat, the customer has to be the one to actually create the power reduction. Even more, a price increase does not even guarantee a shift in peak loads. Summit Blue Consulting oversaw a four-year study in Boulder, Colorado where roughly 1000 customers were given hourly RTP rates for electricity in the summers (Summit Blue 2007). Table 2.1 reports the results of their study where they observed customer cost savings in three of the four years, indicating loads were shifted from the periods of high demand (triggered by the rest of the grid).

Table 2.1: Annual bill Impacts and Weather (Summit Blue 2007)

Year	Average Bill	Average Monthly kWh	Savings	Average Energy Price [¢/kWh]	Maximum Summer Price [¢/kWh]	Cooling Degree Days	Summer Days with High Price Notifications
2003	$51.10	630	20.1%	3.2	12.4	659	9
2004	$56.99	648	11.3%	3.8	12.5	574	7
2005	$77.82	758	-6.3%	5.7	19.1	1087	57
2006	$56.32	679	15.5%	5.1	36.6	937	10

The one exception was 2005 where the pilot customers actually spent more on their monthly bill. This suggests that they used more energy in peak times than the price points were supposed to yield. The cooling degree days were only higher in the summer of 2005 by roughly 10% compared to 2006, but this small difference resulted in over five times as many high price notifications. One can speculate the customers were overwhelmed by the number of notifications and eventually started to ignore

them, incurring the higher usage in the peaks.

Atlantic Electric Company has been testing cycling air conditioners and shutting off hot water heaters using FM radio signals during peak demand on hot afternoons in New Jersey since 1987. Kempton et al. (1992) analyzed the data for several levels of peak duty cycle control including: 5 min off for every 15 min on (25% duty cycle reduction) and 7.5 minutes off for every 7.5 minutes on (50% duty cycle reduction). They were unable to see significant savings in the 25% duty cycle reduction when compared to standard operating air conditioning load on similar days. In this test group 60% of the sample already was cycling at or below the 75% during the peak so no potential was there. Using 50% cycling during the peak event they observed an average of 0.87 kW of load reduction when considering houses with duty cycles above 50% (84% of the sample). If only the top 50% of the sample is considered the load reduction averaged 1.34 kW. Figure 2.8 plots the air conditioning load curves for one customer on a control day and one where the direct control cycling was used.

One metric that is difficult to conclude from the study is how much the reduced cycling of the air conditioners increased the temperature of the homes. Using customer surveys users identifying their experience as "enough cooling" reported an average of 0.18 °C (0.33 °F) indoor temperature increase during these 2 to 5-hour events. A 1.6 °C (3.9 °F) average increase was reported by 25% of the sample which identified as receiving "not enough cooling." Because less than half the sample even reported their temperatures and comfort levels in the surveys and no identification
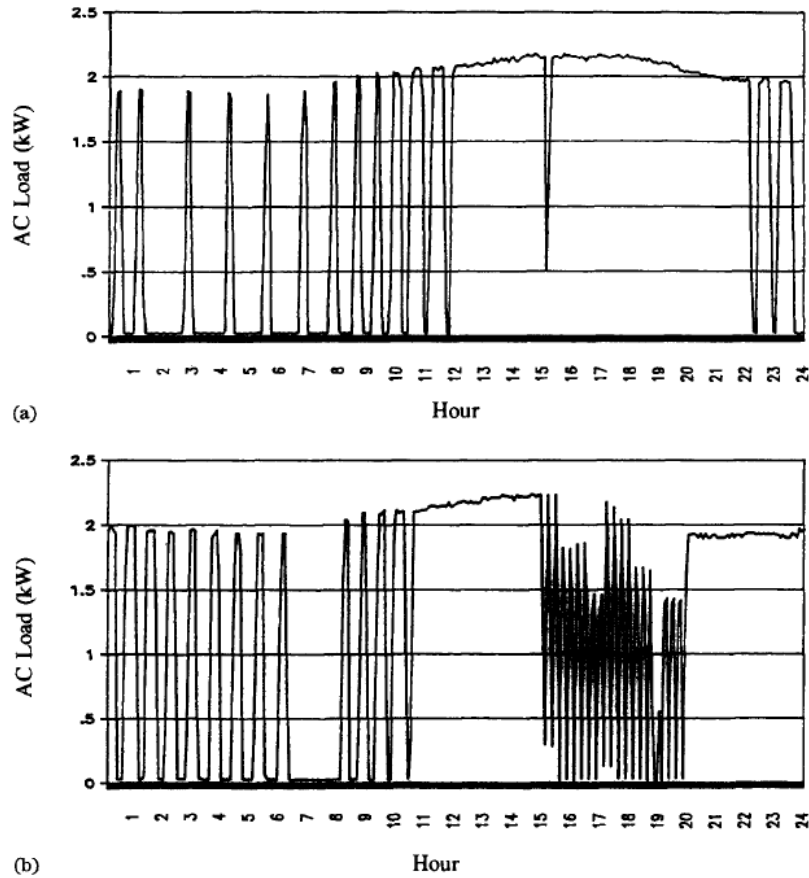
Figure 2.8: Sample 5-min air conditioner load data for one customer for (a) July 13, 1987 without direct load control and (b) July 9, 1987 with 50% load control.

was given for 25% and 50% cycling, it is difficult to determine at what percent cycling would users complain and not want to continue participating.

Controlling the setpoint temperatures of customer thermostats is another direct control approach to reduce peak loads. Pacific Gas and Electric tested thermostat control and direct cycling of air conditioners in the region surrounding Stockton, CA in the late summer of 2007 (KEMA 2008). Direct cycling customers observed peak savings in the range of 0 to 1.34 kW per air conditioning unit using TrueCycle controlled duty cycles. TrueCylce is a variable amount of cycling designed to be half the standard operating duty cycle at those conditions. The thermostats received steep

and gradual setpoint increases from 23.9 °C (75 °F) to a maximum of 26.1 °C (79 °F). This strategy was able to yield a range of 0 to 0.89 kW per air conditioning unit on the different test days. Figure 2.9 plots the projected power reduction per unit along with the actual observed reduction on the test days verses the region's daily average temperature (in °F). The discrepancy between the two treatment groups and the predicted rates show the weakness in only using a few test days to forecast load availability.
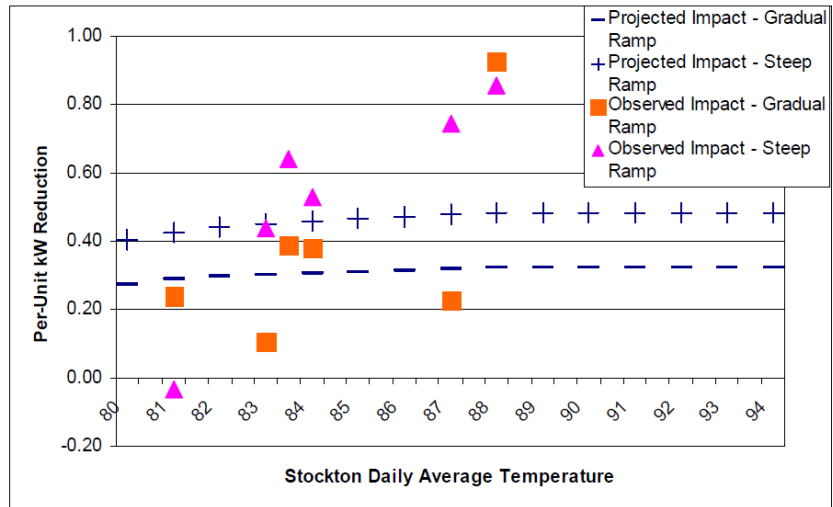
Figure 2.9: Impact projections for thermostat participants with summer 2007 estimated impacts (KEMA 2008).

This and other studies show how direct air conditioning control can remove significant loads but customer responses have not always been positive. Originally reported capacity by utilities has decreased in response to customers opting out of peak load programs or reducing their commitments. A great example of this occurred with Baltimore Gas and Electric's (BGE) PeakRewards program in the summer of 2011(Kay and Baughman 2011). On July 22, 2011, a day with record-breaking heat across the BGE service territory, most of the 450,000 program base was cycled to

20

shift peak loads. Indoor temperatures were reported as high as 37.2 °C (99 °F). It was also reported that customers were so displeased that nearly 10% of the participants called to complain and 1% canceled or reduced their level of participation by the time of the article (July 25, 2011). Also note that these numbers could be much higher because the contact phone lines were at full capacity and the article appeared immediately following the event, before all agitated customers had the chance to change their service.

Newsham and Bowker (2010) presented a collection of seven studies where utilities directly controlled air conditioners during peak periods. Figure 2.10 plots these results.
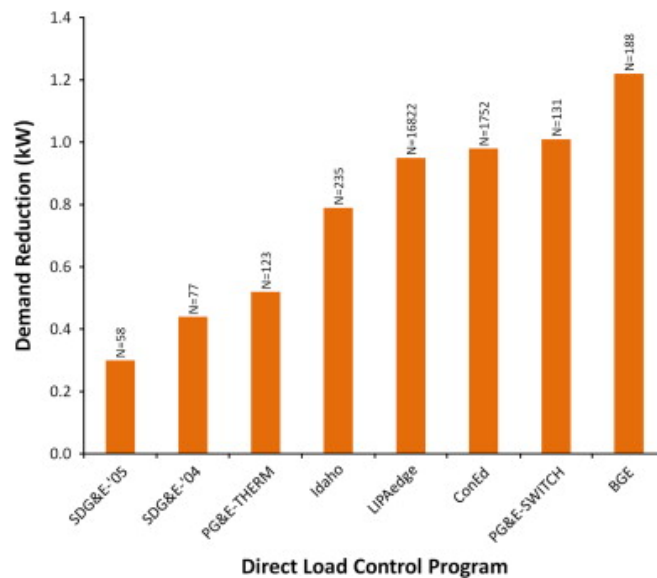
Figure 2.10: Reported average reduction in electricity use per house for the reviewed direct load control studies (Newsham and Bowker 2011).

Figure 2.10 shows the inconsistency between the reported numbers of each study, despite all being air conditioning control. Many variables in each study are different including the geographical location. This impacts the climate and size of air

21

conditioners. The studies also use different strategies to remove the load. Nonetheless, there has to be a more efficient way to forecast how much load is available with direct control demand response using air conditioning than running season long trials.

2.3 Load Forecasting

Independent of demand response, forecasting loads in general is very useful to utilities, energy traders, and power consumers. Several methods were explored to baseline the performance of the energy models presented in the following sections.

Cooling degree days are widely used in load forecasting and weather normalizing energy data. One value is determined per day by averaging the daily high and low outdoor air temperature and subtracting a base where the house should be in balance with the outdoor environment (15.5 °C was used here).

$$CDD = \frac{T_{\infty,high} + T_{\infty,low}}{2} - 15.5 \, [°C] \qquad\qquad (2.1)$$

Traditionally the outdoor air temperature used in these calculations comes from the nearest airport weather station. To be consistent, the same outdoor air temperature measured by the closest WeatherBug weather station will be used for this and all other models. An example of correlating electrical power to the cooling degree days for a house in Houston, TX is shown in Figure 2.11. The hourly average power readings are binned by 0.25 CDD and fit to a linear correlation using ordinary least squares (OLS).
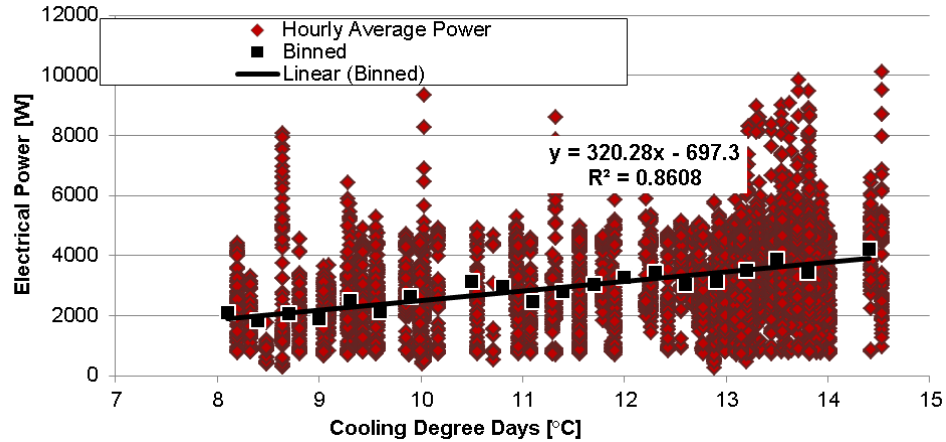
Figure 2.11: Electrical power correlation with CDD for a house in Houston, TX.

Correlations using the cooling degree day method are generated using 250 houses in Houston, TX with data from July 1 to September 30, 2012. These calculations and actual consumption data are summed daily to produce Figure 2.12. This plots use the shading convention of red for a model over-prediction and blue for a model under-prediction. Figure 2.13 plots the hourly predicted and energy consumption data for the week of July 18 to July 25, 2012. This week was chosen to highlight because there was a large range in daily peak consumption.
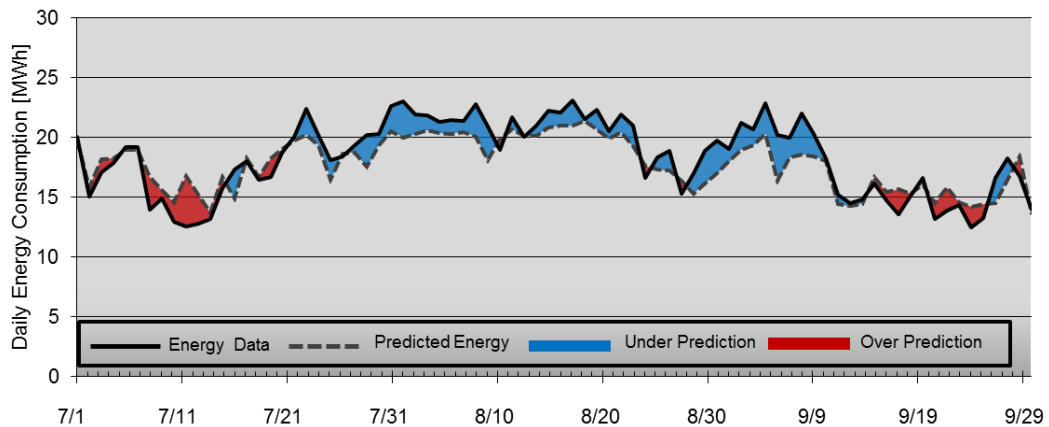
Figure 2.12: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using CDD correlations.

23
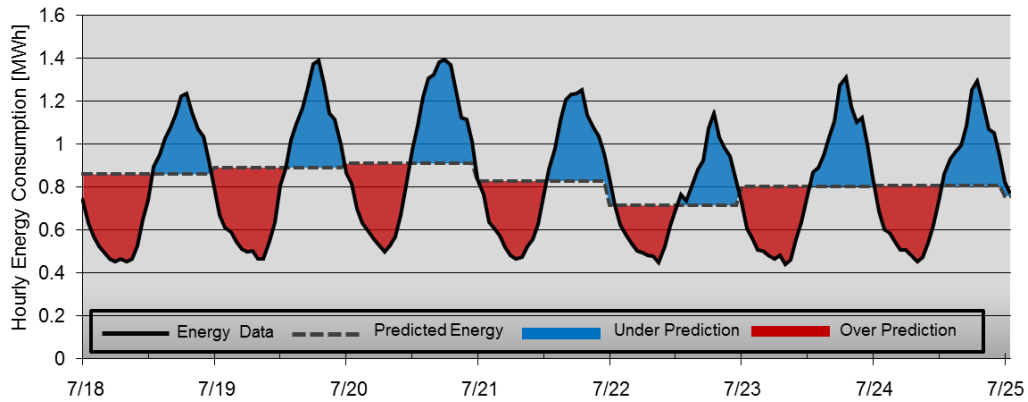
Figure 2.13: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using CDD correlations.

The same predicted energy consumption value is observed throughout a day because only one value can be determined using the CDD technique. If hourly time resolution is required the consumption can be predicted using the average load over the season. Figure 2.14 and Figure 2.15 plot the daily and hourly aggregate energy consumption measured and calculated using the average load method.
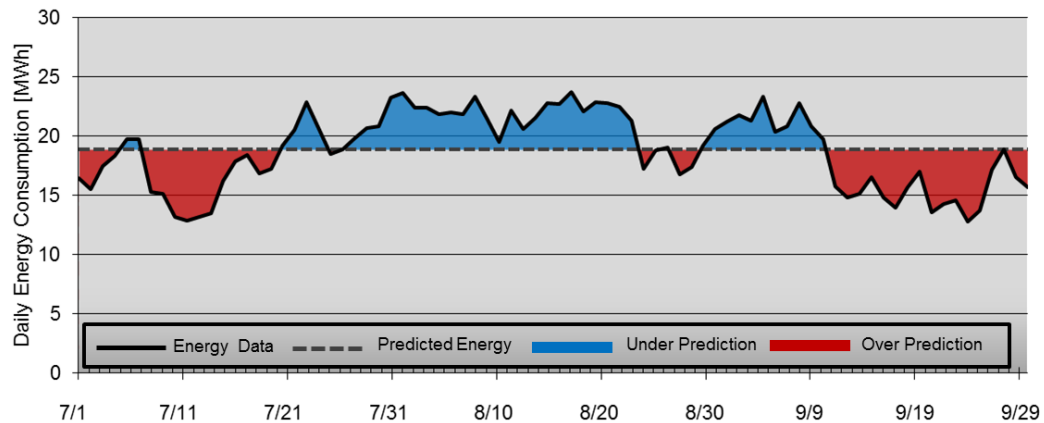
Figure 2.14: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from July 1 to September 30, 2012 calculated using the average hourly load from the entire summer.
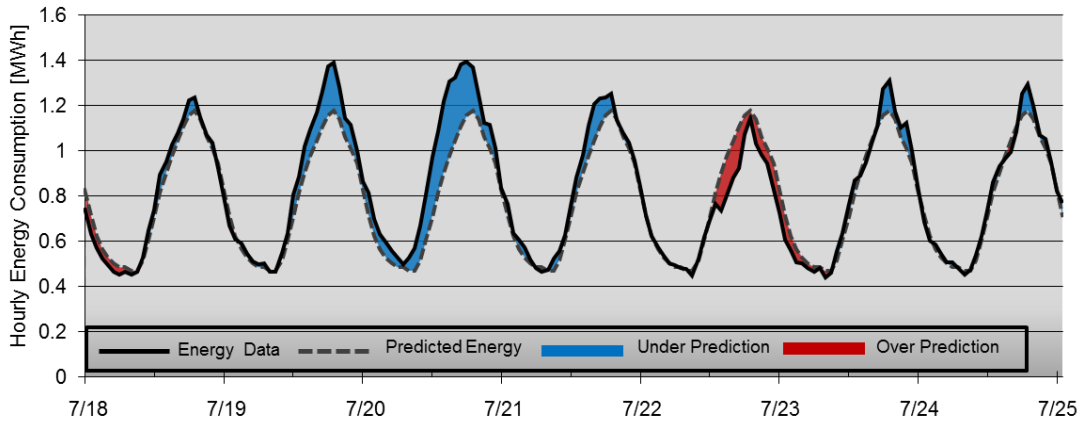
Figure 2.15: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using the average hourly load from the entire summer.

The average load method is unable to predict different daily consumption because the same values are always used. This also makes the hourly predictions poor when days are warmer or cooler than the average. A method designed to improve this is the Middle-Eight-of-Ten-Like-Days method. The Electric Reliability Council of Texas (ERCOT) employs this method in their Emergency Interruptible Load Service program to predict consumption and the baseline load reduction of customers. This method assumes that the energy consumption of recent days would most likely reflect what would occur on the day of an event. Ten days are selected preceding the event that fall into the same category (weekdays or weekend and holiday). The lowest and highest daily energy consumption sum are removed to produce the middle eight days, and a time average of this became the prediction for the current hour. This analysis is performed for each individual house, and aggregated for the entire 250 house set to produce Figure 2.16 and Figure 2.17.

25
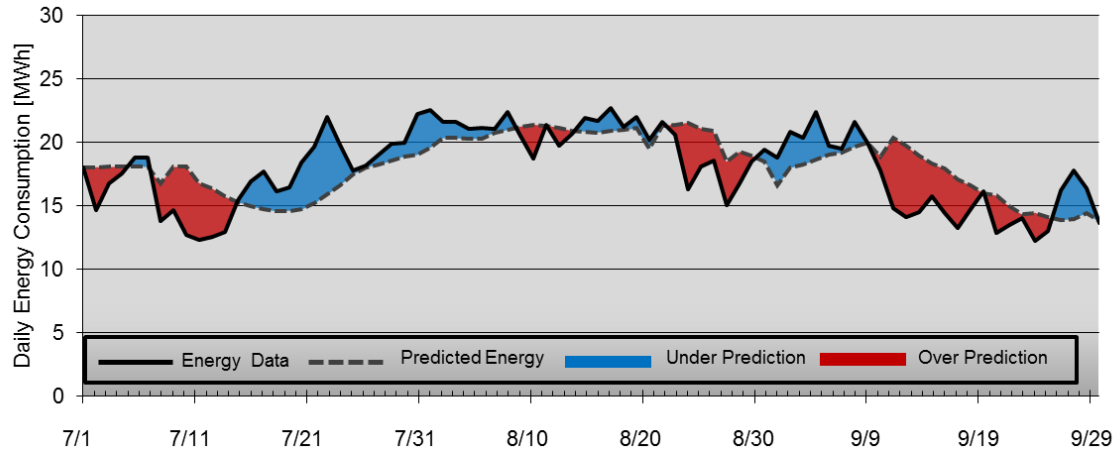
Figure 2.16: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from July 1 to September 30, 2012 calculated using the Middle-Eight-of-Ten-Like-Days method.
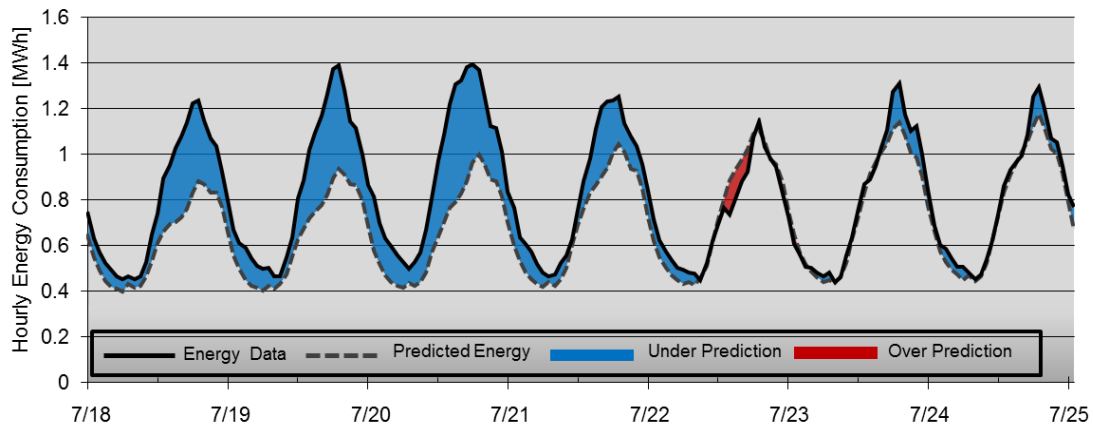
Figure 2.17: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using the Middle-Eight-of-Ten-Like-Days method.

The Middle-Eight-of-Ten-Like-Days method is highly dependent on the previous 10 days. This method does not predict the energy consumption well when the weather changes, and also struggles with the extreme weather that causes grid emergencies.

Directly correlating the indoor temperature measured at a local weather station to the power data is another method used to forecast loads. The power measurements are binned by outdoor temperature and fitted to a linear correlation using OLS. An example correlating the power data to the outdoor temperature for a house in Houston, TX is shown in Figure 2.18. This method is performed for each individual house, and aggregated for the entire set to produce Figure 2.19 and Figure 2.20.

$$y = 223.11x - 3520.8$$
$$R^2 = 0.9694$$
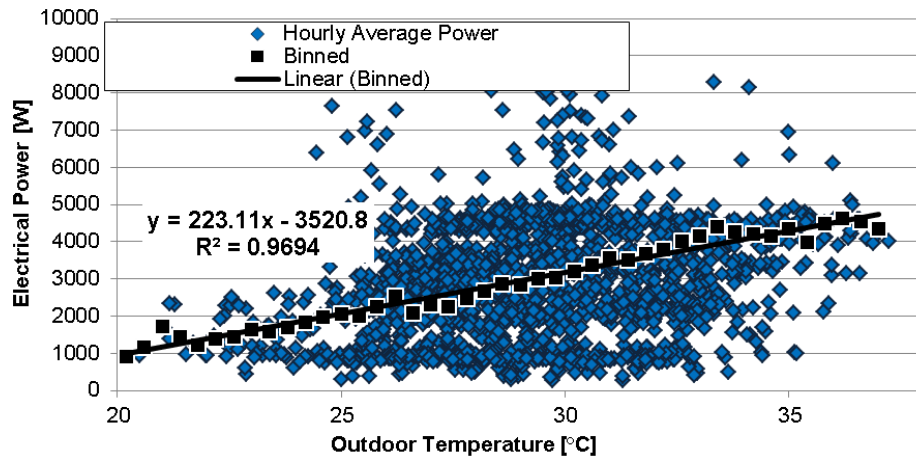
Figure 2.18: Electrical power correlation with outdoor air temperature for a house in Houston, TX.
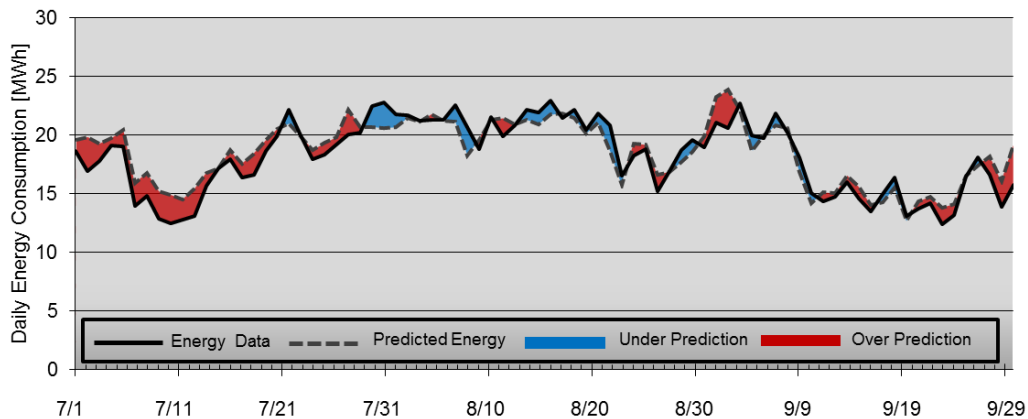
Figure 2.19: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from July 1 to September 30, 2012 calculated using the outdoor temperature correlation method.
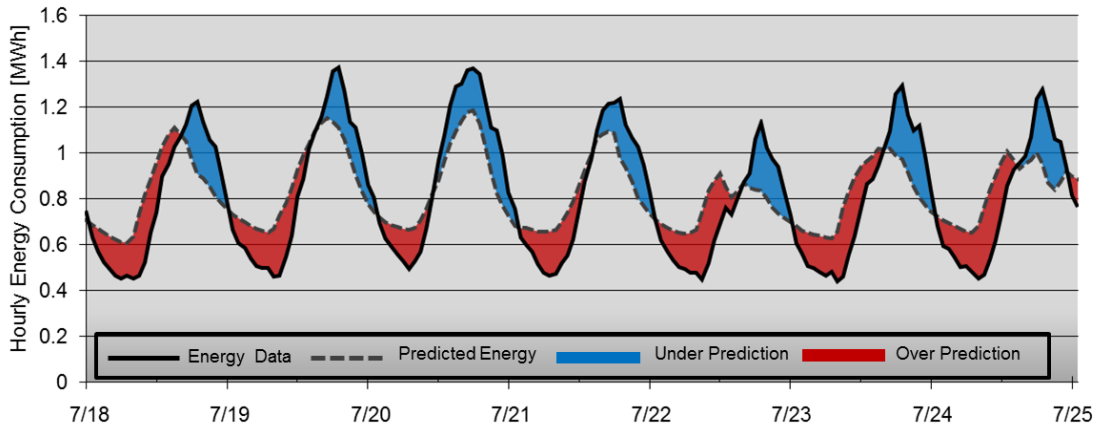
Figure 2.20: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using the outdoor temperature correlation method.

The outdoor temperature correlation method can be improved by delaying the temperature by a time constant. For this set four hours was observed to produce the best fit. Figure 2.21 plots the energy consumption using shifted outdoor temperature.
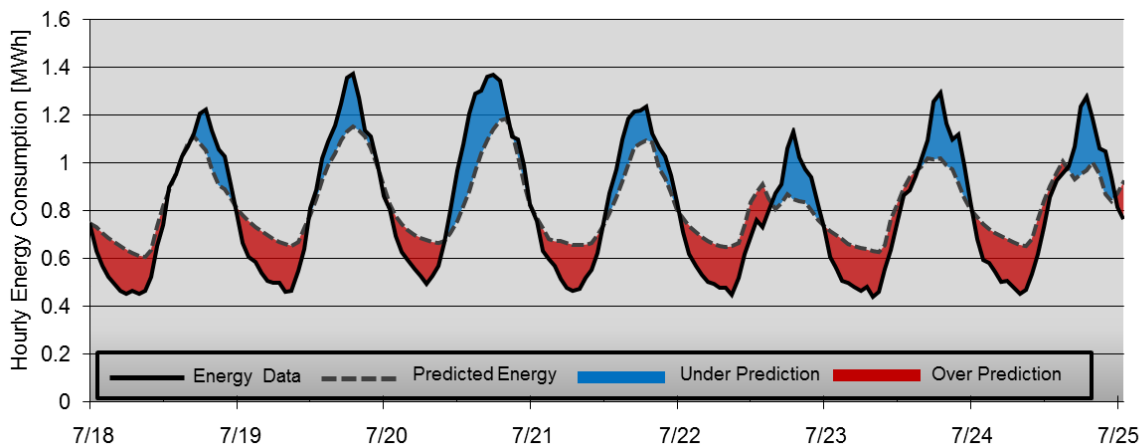
Figure 2.21 Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using the outdoor temperature correlation method shifted by a 4 hour time delay.

Mathieu et al. (2011) used a regression-based baseline model to compute the errors associated with predicted DR removable load. They designed a metric to measure the amount of error that was associated with the baseline model rather than the real variability in the removed load. In the majority of the 38 commercial and industrial buildings they sampled, the discrepancy between the predicted and actual removed load was largely in part due to the baselining method. This is expected when using baseline methodology as poor as ERCOT's middle 8 of 10 method presented with Figure 2.16 and Figure 2.17. Therefore, the forecasting accuracy serves as the primary evaluation metric for the models presented in this dissertation. A summary of the performance of these four prediction methods is given in Table 3.12.

2.4 Motivation

It was determined by reviewing the current literature and trials performed by large electricity utilities that an efficient method to accurately predict how much peak load is removable by directly controlling thermostats does not exist. A large discrepancy in the reported capacities of the available studies and pilot experiments was observed. The main root of this cause is the variability of the studies and programs. Different control strategies were employed to adjust the thermostats, and each study was only conducted in one location of the country. Now, the only option to get an accurate representation of how much power is available is to perform a pilot study in the targeted service territory. Not only does this take time, it also can be expensive. This dissertation proposes analyzing the data from internet-connected thermostats, smart power meters, and local weather data to estimate the available removable load given any operating control strategy and location.

The thermostat and smart meter data is combined with local weather observations and forecasts to generate unique building energy models for each connected thermostat. After reviewing the literature on effective models, the thermodynamic energy equations were modeled as a grey-box system with effective parameter coefficients. These coefficients are determined using a Genetic Algorithm optimization solver developed to minimize the RMS error between the modeled indoor air temperature and the actual data. This system was designed to be both simple and computationally light so simulations can be made quickly and without requiring any further information from the actual house. This model is referred to as the thermostat building energy model in this dissertation, and was designed to:

1. Predict the runtime and energy consumption of houses under standard operation conditions

2. Execute demand response load removal events by controlling thermostats in a trial with utilities

3. Provide information to improve the current techniques of thermostatically controlled demand response so customers remain in control and avoid mass exoduses from demand response programs

4. Forecast the available removable load for any location, weather condition, and control strategy

A second building energy model was also developed to forecast loads when internet-connected thermostat data was not available. This model is referred to as the smart meter only building energy model and was built using smart power meter and local weather station data. The derivation and study of the performance of each model

is presented in the following Building Energy Model Chapter. Also note a glossary of terms is found at the end of the dissertation.

Chapter 3: Building Energy Models

3.1 Introduction

Two models were developed to forecast the energy consumption in houses. The first is a thermodynamic model that is trained using internet-connected thermostats, weather network, and smart meter data and referred to as the Thermostat Building Energy Model. The second only used the weather network and smart meter data and is referred to as the Smart Meter Only Building Energy Model. The derivation and specifics of each model is described in the following sections. The performance of each is summarized in the Energy Prediction Summary section.

3.2 Thermostat Building Energy Model

3.2.1 Model Derivation

The single largest consumer of energy in a home is usually the HVAC system. Because the HVAC energy consumed is driven by the indoor air temperature measured by the thermostat, modeling the temperature is the logical way to best predict the energy consumption. In the grey-box Thermostat Building Energy Model this was done by identifying the important energy flows within a home. It was assumed that the interior of the home was separated from the outdoor environment by a one-dimensional wall. Heat transfer to the outer surface of this wall occurs through convection with the outside air and through incident solar radiation. Heat diffuses through the wall to the inside surface, which then transfers heat to the indoor air through convection. The indoor air is represented by a single node to which heat is

transferred via convection on the inside wall, solar radiation passing through the windows, air infiltration through gaps and cracks in the envelope, and convection with the internal thermal mass (e.g., floors, furniture, interior walls). The HVAC system adds or removes heat when called upon to regulate the thermal comfort. These energy flows are represented using the thermoelectricity analogy in Figure 3.1 where the Ts, Rs, Cs, and Qs represent temperatures, thermal resistances, capacitances, and heat flows, respectively. The subscripts ∞, w, d, i, a, and m denote the ambient, wall, windows, wall grid index, indoor air, and internal mass, respectively.
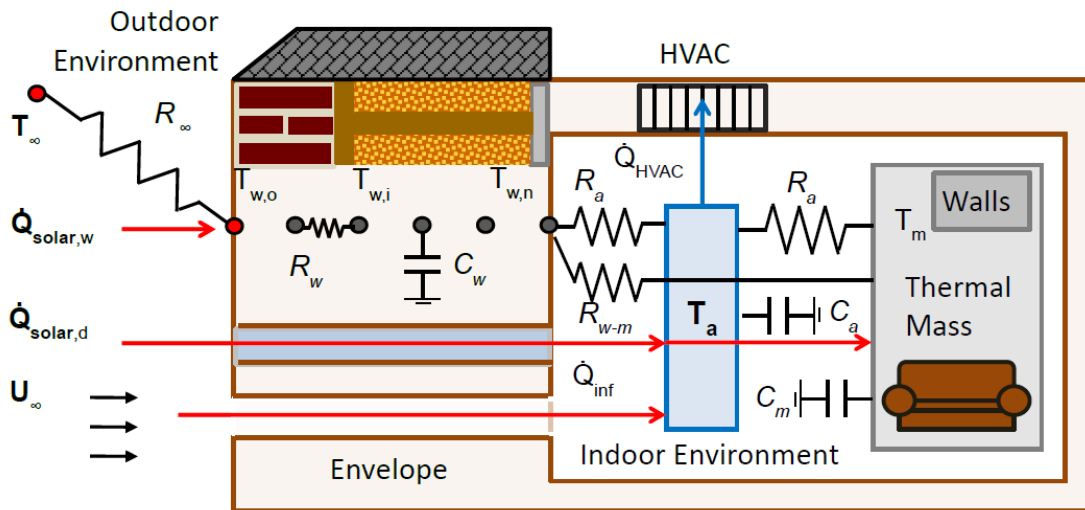
Figure 3.1: Thermostat building energy model thermoelectric analogy model. Note the system input data set is bold (Outdoor air temperature, solar load, wind speed, indoor air temperature, and HVAC status).

The thermal resistance between the outside ambient air and the wall of the house is comprised of both natural and forced convection. The natural convection on the outside wall is modeled as a vertical plate using the Churchill and Chu correlation (Incropera 1990):

$$\text{Nu}_{natural,\infty} = \left(0.825 + \frac{0.387 \, \text{Ra}^{1/6}}{1+\left(\frac{0.491}{\text{Pr}^{9/16}}\right)^{8/27}}\right)^2 \tag{3.1}$$

Where the wall average Nusselt Number (Nu), Rayleigh Number (Ra), and Prandtl Number (Pr) are defined as:

$$\text{Nu} = \frac{h \, A_w}{k_\infty} \tag{3.2}$$

$$\text{Ra} = \frac{g \, \beta_\infty}{\nu_\infty \, \alpha_\infty}\left(T_{w,0} - T_\infty\right)A_w^3 \tag{3.3}$$

$$\text{Pr} = \frac{\nu_\infty}{\alpha_\infty} \tag{3.4}$$

The forced convection off the outside wall is modeled using the turbulent boundary layer equation:

$$\text{Nu}_{forced,\infty} = \left(0.296 \, \text{Re}^{4/5} \, \text{Pr}^{1/3}\right) \tag{3.5}$$

Where the wall average Reynolds Number (Re) is defined as:

$$\text{Re} = \frac{U_\infty A_w}{\nu_\infty} \tag{3.6}$$

The Archimedes number (Ar) is used to determine if including both convections terms is valid. This is true if the ratio is ≈ 1.

$$\text{Ar} = \frac{\text{Gr}}{\text{Re}^2} = g \, \beta_\infty U_\infty^2 (T_w - T_\infty) A_w \tag{3.7}$$

Values $\gg 1$ are observed when the solar loading heats up the outside surface of the wall and values $\ll 1$ are observed with higher wind speeds so both convection regimes are combined. The composite mixed convection coefficient is modeled using:

$$h_{mixed,\infty} = \frac{k_\infty}{A_w}\left(\text{Nu}_{natural,\infty}^3 + \text{Nu}_{forced,\infty}^3\right)^{1/3} \tag{3.8}$$

The total outdoor wall heat transfer is:

$$\dot{Q}_{conv,\infty} = h_{mixed,\infty} \, A_w \frac{dT}{dt} \tag{3.9}$$

Solar radiation is incident on the outer wall of the structure with a fraction of that passing directly through the windows into the indoor environment. The ratio of the window area to total wall area is the fenestration ratio. The two solar heat fluxes are modeled using this ratio as:

$$\dot{Q}_{solar,w} = \dot{Q}_{solar,in} \, (fr - 1) \tag{3.10}$$

$$\dot{Q}_{solar,d} = \dot{Q}_{solar,in} \, fr \tag{3.11}$$

$$fr = \left(\frac{A_d}{A_w}\right) \tag{3.12}$$

Where the "in" subscript represents total incident solar radiation and "fr" is the fenestration ratio. The infiltration of outdoor air is driven by the pressure and temperature difference between the indoor and outdoor environments. All the cracks and gaps in the building envelope are modeled as a single pipe. The wind increases the pressure at the pipe inlet and outdoor air flows in. The heat transfer into the indoor environment from infiltration is modeled as:

$$\dot{Q}_{inf} = \dot{m}cp_a(T_\infty - T_a) \tag{3.13}$$

Figure 3.2 displays a schematic of the infiltration model.
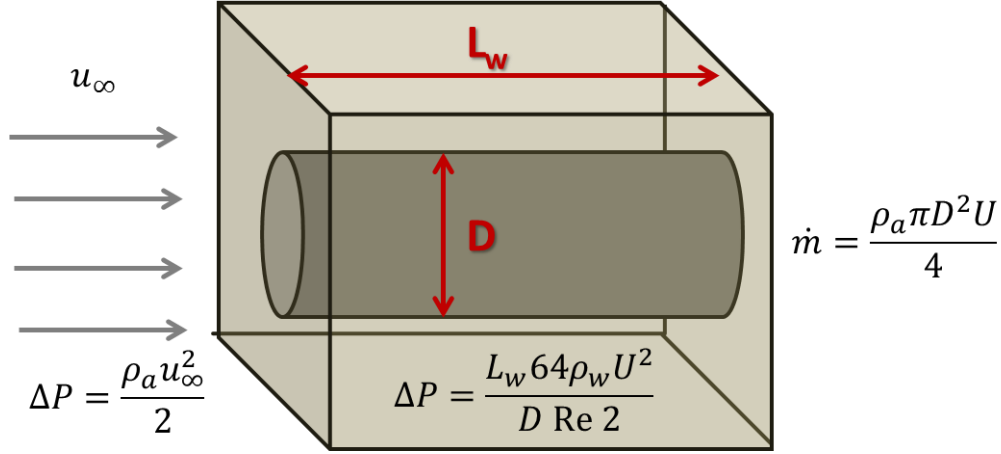
Figure 3.2: Infiltration modeling schematic.

The pressure difference at the inlet of the pipe is given in Eq. 2.14 and the pressure drop through the pipe assuming laminar flow is Eq. 2.15 with 64/Re representing the friction factor and U the average velocity.

$$\Delta P = \frac{\rho_a u_\infty^2}{2} \tag{3.14}$$

$$\Delta P = \frac{L_w 64 \rho_w U^2}{D \, \text{Re} \, 2} \tag{3.15}$$

The mass flow rate through the pipe induced form this pressure difference is:

$$\dot{m} = \frac{\rho_a \pi D^2 U}{4} \tag{3.16}$$

The system of the two pressure difference equations are solved for the mass flow rate resulting in:

$$\dot{m} = \frac{\rho_a^2 \pi D^4}{256 L_w \mu_w} u_\infty^2 \tag{3.17}$$

The diffusion though the wall includes both conduction and energy storage. In one dimension the diffusion equation is written as:

$$\frac{\partial \varnothing}{\partial t} - \alpha_w \frac{\partial^2 \varnothing}{\partial x^2} = 0 \tag{3.18}$$

36

Because the temperatures differences between the inside wall and indoor air are relatively small, the convection and radiation heat transfer are represented by a single thermal resistance. This is also the case for the air to the thermal mass and inside wall to thermal mass heat transfer:

$$\dot{Q}_{w-a} = \frac{R_a(T_a - T_{w,n})}{A_w} \qquad (3.19)$$

$$\dot{Q}_{a-m} = \frac{R_a(T_m - T_a)}{A_m} \qquad (3.20)$$

$$\dot{Q}_{w-m} = \frac{R_{w-m}(T_m - T_w)}{A_w} \qquad (3.21)$$

The air also has its own thermal mass that must be overcome before a temperature change can occur.

$$\dot{Q}_{a,capacity} = V_a \rho_a c p_a \frac{dT_a}{dt} \qquad (3.22)$$

The HVAC system is modeled to input or remove heat at a rate dependent on the outdoor ambient temperature (when applicable: AC and heat pumps). HVAC manufacture's data sheets report linear correlations between the cool/heating capacity and the outdoor condensing/evaporating air temperatures. In the following equation the slope and intercept of such a correlation are represented by "M" and "B" respectively:

$$\dot{Q}_{HVAC} = M_{capacity}T_\infty + B_{capacity} \qquad (3.23)$$

The thermal capacity of the large, solid objects in a house not associated with the envelope are modeled as:

$$\dot{Q}_{m,capacity} = V_m \rho_m c p_m \frac{dT_m}{dt} \qquad (3.24)$$

Latent loads were left out of the model because the thermostats used neither measure nor control to the indoor relative humidity. The addition of moisture to the indoor environment could be modeled similarly to infiltration in the future. Also internal heat generation from occupants, appliances, and devices were not a part of this specific model because it was deemed too irregular to forecast.

The system is mapped to a grid of 22 points: 20 (n) points to represent the one-dimensional, uniform property envelope wall, one for the indoor air, and the last for the indoor thermal mass. The system of equations is solved explicitly using a finite difference method so the equations are modeled at a future time and the second derivatives in the diffusion equations can be represented numerically using a truncated Taylor-series expansion about the grid point i.

$$\left(\frac{\partial^2 \emptyset}{\partial x^2}\right)_i = \frac{\emptyset_{i-1} + \emptyset_{i+1} - 2\emptyset_i}{\Delta x^2} \tag{3.25}$$

The full system is represented by the following discretized, explicit equations. Subscripts denote the position in the discrete equation grid, and the superscripts with "t" or "t+1" represent the current or future time steps respectively. The outdoor wall node (i = 0) with R_∞ simplifying all the outdoor convection terms:

$$R_\infty A_w L_w \left(T_\infty^{t+1} - T_{w,0}^{t+1}\right) + k_w A_w \frac{T_{w,1}^{t+1} - T_{w,0}^{t+1}}{\Delta x} + \dot{Q}_{solar,in}(\text{fr} - 1)A_w L_w =$$

$$\Delta x A_w \rho_w c p_w \frac{T_{w,0}^{t+1} - T_{w,0}^t}{2\Delta t} \tag{3.26}$$

The 18 interior wall nodes (i = 1 to n-1 or 19):

$$k_w A_w \frac{T_{w,i-1}^{t+1} - T_{w,i}^{t+1}}{\Delta x} + k_w A_w \frac{T_{w,i+1}^{t+1} - T_{w,i}^{t+1}}{\Delta x} = \Delta x A_w \rho_w c p_w \frac{T_{w,i}^{t+1} - T_{w,i}^t}{\Delta t} \tag{3.27}$$

The indoor wall node (i = n):

$$k_w A_w \frac{T_{wn-1}^{t+1} - T_{w,n}^{t+1}}{\Delta x} + R_a A_w \left(T_a^{t+1} - T_{w,n}^{t+1}\right) + R_{w-m} A_w \left(T_m^{t+1} - T_{w,n}^{t+1}\right) =$$

$$\Delta x A_w \rho_w c p_w \frac{T_{w,n}^{t+1} - T_{w,n}^t}{2\Delta t} \tag{3.28}$$

The indoor air node (a):

$$R_a A_w (T_n^{t+1} - T_a^{t+1}) + R_a A_m (T_m^{t+1} - T_a^{t+1}) + \dot{Q}_{solar,in}^{t+1}(\text{fr}) +$$

$$\dot{m} c p_\infty u_\infty^{t+1^2} (T_\infty^{t+1} - T_a^{t+1}) + HVAC_{mode}^{t+1} (M_{capacity} T_\infty + B_{capacity}) =$$

$$V_a \rho_a c p_a \frac{T_a^{t+1} - T_a^t}{\Delta t} \tag{3.29}$$

The indoor thermal mass node (m):

$$R_a A_m (T_a^{t+1} - T_m^{t+1}) + R_{w-m} A_m (T_n^{t+1} - T_m^{t+1}) + \dot{Q}_{solar,in}^{t+1}(\text{fr}) =$$

$$V_m \rho_m c p_m \frac{T_m^{t+1} - T_m^t}{\Delta t} \tag{3.30}$$

The equations are then non-dimentionalized with respect to length and time using the following non-dimensional length and time scales:

$$\Delta x^* = l_w \Delta x, \tag{3.31}$$

$$\Delta t^* = \alpha_w \frac{\Delta t}{l_w^2} \tag{3.32}$$

It was determined all the building parameters in the discretized heat transfer equations could be represented as 12 coefficients. Table 3.1 shows the properties represented by the 12 coefficients followed by the non-dimensional heat transfer equations.

Table 3.1: Effective building parameter heat transfer equation coefficients

Coefficient	Proportional To	Equation
1	Wall Thermal Mass	$c_1 = A_w \rho_w cp_w$
2	Wall Conductivity	$c_2 = k_w l_w$
3	Infiltration Factor	$c_3 = \rho_a cp_a \dfrac{\rho_a^2 \pi D^4}{256 L_w \mu_w}$
4	Indoor Air Thermal Resistance	$c_4 = R_a A_w$
5	Wall Sun Factor	$c_5 = (\text{fr} - 1)A_w$
6	Window Sun Factor	$c_6 = (\text{fr})A_d$
7	HVAC Capacity Equation Intercept	$c_7 = M_{capacity}$
8	Wall to Mass Thermal Resistance	$c_8 = R_{w-m} A_m$
9	Outdoor Forced Convection	$c_9 = fun(A_w, \alpha_\infty, \nu_\infty, k_\infty)$
10	Outdoor Natural Convection	$c_{10} = fun(A_w, \alpha_\infty, \nu_\infty, k_\infty, g, \beta_\infty)$
11	Internal Thermal Mass	$c_{11} = V_m \rho_m cp_m$
12	Indoor Air Thermal Mass	$c_{12} = V_a \rho_a cp_a$

The following non-dimentional heat transfer equations are used to represent the building. The data sources are bolded to help clarify what is known. The 12 coefficients are the independent variables adjusted to determine the indoor air temperature (dependent variable). The equation for the outdoor wall node ($i = 0$) is:

$$\frac{R_\infty(\boldsymbol{T_\infty^{t+1}} - T_{w,0}^{t+1})}{c_2} + \frac{T_{w,1}^{t+1} - T_{w,0}^{t+1}}{\Delta x^*} + \frac{c_5 \boldsymbol{\dot{Q}_{solar,in}}}{c_2} = \frac{\Delta x^*(T_{w,0}^{t+1} - T_{w,0}^t)}{2\Delta t^*}$$

$$\text{where: } R_\infty \text{ is a } Function(c_9, c_{10}, \boldsymbol{T_\infty}, \boldsymbol{U_\infty}) \tag{3.33}$$

The 18 interior wall node ($i = 1$ to 19) equations are:

$$\frac{T_{w,i-1}^{t+1} - T_{w,i}^{t+1}}{\Delta x^*} + \frac{T_{w,i+1}^{t+1} - T_{w,i}^{t+1}}{\Delta x^*} = \frac{\Delta x^*(T_{w,i}^{t+1} - T_{w,i}^t)}{\Delta t^*} \tag{3.34}$$

The indoor wall node ($i = n$) equation is:

$$\frac{T_{wn-1}^{t+1} - T_{w,n}^{t+1}}{\Delta x^*} + \frac{c_4(T_a^{t+1} - T_{w,n}^{t+1})}{c_2} + \frac{c_8(T_m^{t+1} - T_{w,n}^{t+1})}{c_2} = \frac{\Delta x^*(T_{w,n}^{t+1} - T_{w,n}^t)}{2\Delta t^*} \tag{3.35}$$

The indoor air node (a) equation is:

$$c_4(T_n^{t+1} - T_a^{t+1}) + c_4(T_m^{t+1} - T_a^{t+1}) + c_6\dot{Q}_{solar,in}^{t+1} + c_3 u_\infty^{t+1}{}^2(T_\infty^{t+1} - T_a^{t+1}) +$$

$$(c_7 - c_8 \boldsymbol{T}_\infty^{t+1})\boldsymbol{HVAC}_{mode}^{t+1} = \frac{c_{12}(T_a^{t+1} - T_a^t)}{\Delta t^*} \tag{3.36}$$

The internal thermal mass node (m) equation is:

$$c_4(T_a^{t+1} - T_m^{t+1}) + c_8(T_n^{t+1} - T_m^{t+1}) + c_6\dot{Q}_{solar,in}^{t+1} = \frac{c_{11}(T_m^{t+1} - T_m^t)}{\Delta t^*} \tag{3.37}$$

The coefficients are also normalized by factors of the following magnitudes. The exact factors were not reported to protect the security of the homeowner in the event this model was ever reproduced exactly.

Table 3.2: Normalizing factors for the coefficients

Coefficient	Normalizing Factor
1	100
2	0.0001
3	0.001
4	1
5	10
6	0.001
7	10
8	1
9	1
10	1
11	10
12	10

The discretized non-dimensional heat transfer equations representing the temperatures within the wall are solved explicitly using a tridiagonal matrix algorithm (TDMA) operating in 5-minute time steps (Patankar 1980). The TDMA numerical technique is unable to resolve the power relations in the outdoor convection terms. Because the conditions outside do not fluctuate drastically on this timescale, an approximation is performed using previous time step information.

The previous 6 hours of indoor and outdoor data are used during the "build-up"

phase to establish the temperature profile in the wall and the starting temperature for the internal thermal mass. This is done by setting the indoor air temperature to be equal to the data. Once the time step equals the time at the start of the simulation, the indoor temperature restriction is removed and the value is calculated solely using the model. During the training phase for the model, explained in the section below, the model knows when the HVAC is ON ($HVAC_{mode}$ in the equations). This is in contrast to the control simulations used to predict when the HVAC would be turned ON. The simulation is performed one time step at a time where the HVAC status of the subsequent step is determined using the current model air temperature, setpoint, hysteresis, and HVAC status. Figure 3.3 is a flow chart showing the logic of the control algorithm used in the simulations. It was designed to replicate the logic on the thermostats and is deterministic.
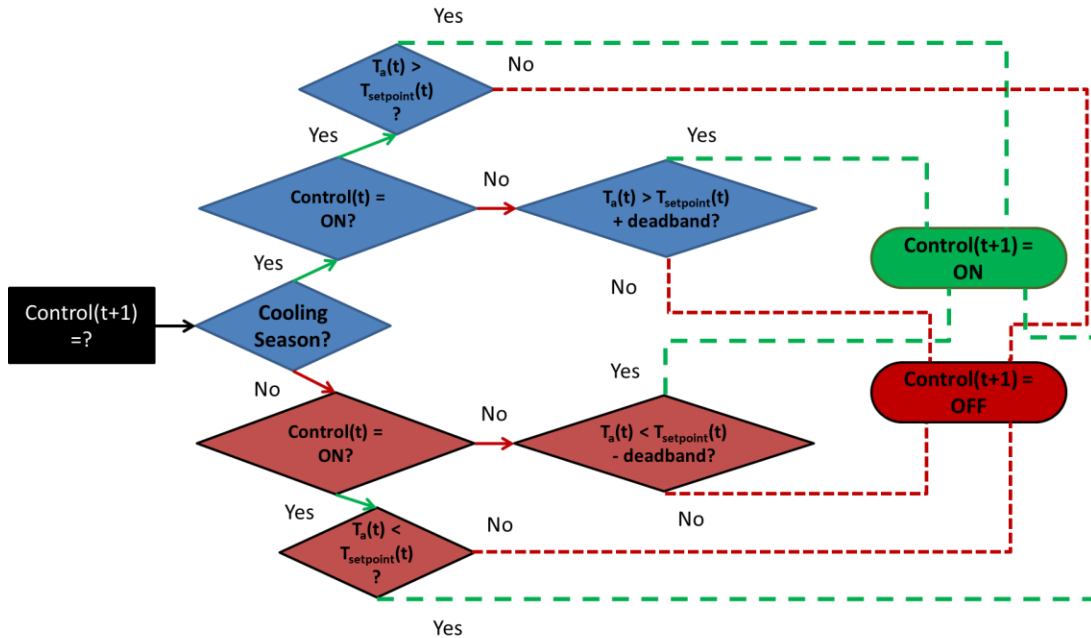
Figure 3.3: Flow chart for the HVAC control algorithm used in the control simulations.

42

Radio Thermostat CT-30, CT-50, and CT-80 internet-connected thermostats were used to obtain the indoor temperature and HVAC status (Control 2013). Figure 3.4 shows the touchscreen display of the CT-80.

Figure 3.4: Radio Thermostat CT-80 internet-connected thermostat.

These thermostats are connected to the internet via a homeowner's Wi-Fi network, and send data to an external server. Users can program and update the thermostat schedule directly on the thermostat, on a website, or model phone / tablet application. The thermostats are accurate to within 0.56 °C (1 °F) at a resolution of 0.28 °C (0.5 °F), and can be programed with seven daily setpoint schedules. They also have a user defined hysteresis which is determined for the energy model by observing the data. The data is recorded in 5-minute intervals as well as when a change is registered. Interpolating between these points allows time steps to have percentages of ON time (2-minutes ON in a 5- minute interval would yield 0.4 ON time). It should be noted the price of these thermostats range from $75-200, and that they require a dedicated 24VAC power source (C-wire).

The weather data used is provided by matching the house location to the closest

Earth Networks WeatherBug weather station, which was usually in the same zip code for this study. These data are either gathered in 5-minute intervals or interpolated to this level from hourly observations. The weather network measurements of temperature are made using crystal oscillator thermometers, wind speed using anemometers, and incident irradiance from pyranometers. An example of the data from one weather station at the different resolutions is shown in Figure 3.5. Both datasets are prone to have missing or erroneous data.
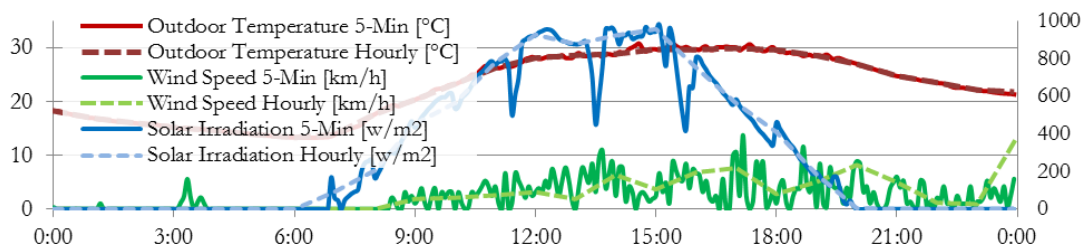
Figure 3.5: Weather observation data at 5-minute and hourly resolutions.

Data was interpolated to fill missing weather using known points when acceptable. Filters were put in place to eliminate unacceptable days from the training and prediction sets, but they were far from foolproof. Noisy data is another nuisance this modeling technique needs to tolerate (e.g., simply walking past a thermostat or opening a door can change the measured temperature by 1°C).

3.2.2 Model Training Phase

Because the model operates as a grey-box system and no physical information about the structure is known, the heat transfer equation coefficients are determined from a training phase. It is shown in 2.2.4 that the ten most extreme days in a month are sufficient to train the system. This method reduces the computational effort three-fold. An extreme day in the set is defined by having the maximum and minimum of

44

the following variables: outdoor temperature, outdoor and indoor temperature range, daily average solar load, and infiltration potential. Best visualized on a spider plot, these 10 "spider days" encompass all the weather and indoor conditions observed in the initial set. The observations are normalized so that the highest value is plotted as a 1 and the lowest as a 0.
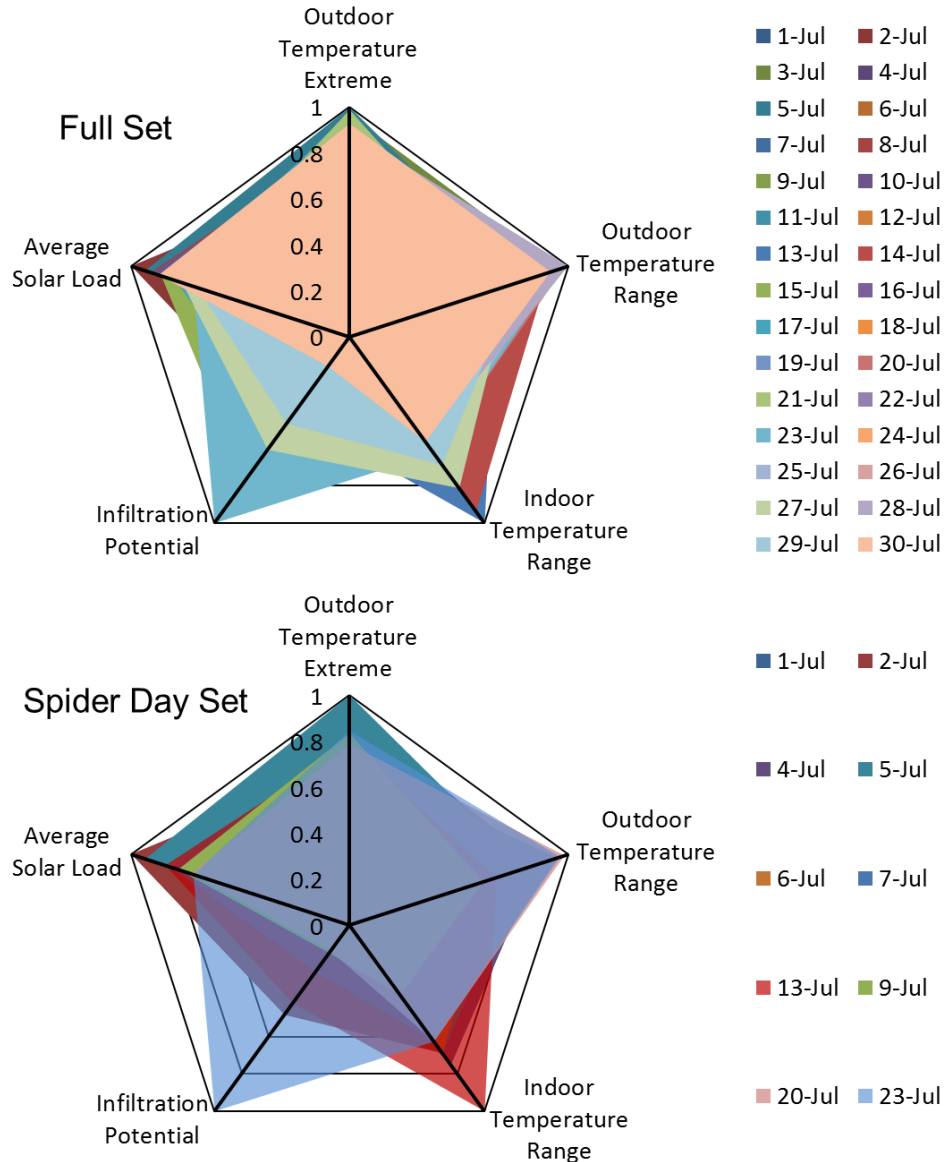
Figure 3.6: Spider plot of the weather and indoor conditions of the full, month-long set (top) and the 10 most extreme "spider days" (bottom).

The values of the coefficients that best define the house are ones that produce the smallest difference between the indoor air temperature data and generated from the model. A root-mean-square (RMS) error between the two temperatures is calculated for the entire 10-day training data (288 points in a day using 5-minute intervals). The coefficient variables are bounded to be greater than or equal to zero to prevent models that do not make physical sense (negative thermal mass, convection, etc.). The problem is formulated as follows:

$$\underset{c_1,c_2,...,c_{12}}{\text{minimize}} \sum_{day=0}^{10} \sum_{t=0}^{288} \left(\left(T_a^t \left(c_1,c_2,...,c_{12} \right) \right)^2 - \left(T_{tstat}^{\ t} \right)^2 \right)^{1/2}$$

$$s.t.:$$

$$c_1,c_2,...,c_{12} \geq 0$$

(3.38)

In this formulation the 12 building parameter coefficients are represented as c_1, c_2, ..., c_{12}. The 10 days evaluated are the spider day subset of the total 30 day set, and the 288 daily time points are for each 5 minutes. The modeled indoor air temperature is T_a and the indoor air temperature data measured by the thermostat is T_{tstat}.

In an effort to quickly produce a modeling system that interfaced with the WeatherBug weather data network and operation tools a custom built Genetic Algorithm (GA) function in visual c# 2012 programing language was originally implemented to generate coefficient models. The GA function was initially tuned to the parameters in Table 3.3: Genetic Algorithm solution parameters to generate acceptable models. These parameters could be optimized for better performance if GA was to continue to be the solution technique (see Tan and Li 2002 for comments on using evolutionary computing in grey-box modeling).

Table 3.3: Genetic Algorithm solution parameters

Parameter	Value
Mutation Rate	0.7
Crossover Rate	0.4
Population Size	25
Generation Size	100
Chromosome Length	12 (coefficients)

GA may not be the most efficient solution technique to determine the building model coefficients. It may also be obtaining coefficient values that are not near an optimum. A simpler approach such as ordinary least squares may be a better technique for determining the coefficient values. Another method was not explored in main model evaluation section of this dissertation because all the applications described in Chapter 3 (and many others not mentioned in this dissertation) interface with the model function formatted this way. However, a subset of 10 houses was tested using MATLAB's GA, Fminunc, Fminsearch, and Lsqnonlin solution functions (2010). A penalty was introduced to bound the coefficient values from being negative when the function did not operate with bounds. The results of this analysis are presented at the end of Section 3.2.4. When the solution system performance needs to be improved it can be retooled to use a simpler technique.

Once known, these coefficients are stored in a database and called upon for simulations. The physical values they represent change with time (e.g., to account for seasonal changes in shading, degradation of AC performance, etc.) so this training phase needs to be repeated periodically.

3.2.3 Model Evaluation Methods

The performance of the presented Thermostat Building Energy Model was tested by investigating how well it could match and predict the indoor air temperature and

thermostat ON time over various conditions. The RMS temperature error from the model coefficient generation is used to evaluate the performance of matching the indoor temperature. However, this error is converted to an average and calculated from the entire 30 day set (not just the spider days). The following equation shows how the temperature error metric is calculated:

$$\text{Temperature Error} = \frac{1}{30 \times 288} \sum_{day=0}^{30} \sum_{t=0}^{288} \left(\left(T_a^t \right)^2 - \left(T_{tstat}^t \right)^2 \right)^{\frac{1}{2}} [°C] \tag{3.39}$$

The predicted ON time metric is calculated by using the control simulation (Figure 3.3) with thermostat setpoints and a control algorithm in place of knowing when the system was actually on. This metric is generated after the coefficient model has been established using the GA, and it also is calculated over the entire 30 day set. A percent error in the ON time is determined using the sum of the predicted and actual ON times. The following equations show how the ON time error metric is calculated.

$$\text{ON Time Data Sum} = 5 \times \left(\sum_{day=0}^{30} \sum_{t=0}^{288} HVAC_{mode}^t \right) [\text{minutes}] \tag{3.40}$$

$$\text{Predicted ON Time Sum} = 5 \times \left(\sum_{day=0}^{30} \sum_{t=0}^{288} Control(t) \right) [\text{minutes}] \tag{3.41}$$

$$\text{ON Time Error} = \frac{(\text{Predicted ON Time Sum}) - (\text{ON Time Data Sum})}{(\text{ON Time Data Sum})} \tag{3.42}$$

When comparing groups of multiple house models these meterics are combined to form group average values. The average daily model predicted ON time sum is the average predicted ON time sum on the average day for the entire group:

Average Daily Model Predicted ON Time Sum =

$$\frac{1}{\#tstats} \times \frac{1}{30}\left(\sum_{tstat=1}^{\#tstats} \text{Predicted ON Time Sum}(tstat)\right)[\text{minutes}] \quad (3.43)$$

The average daily ON time data sum follows the same logic:

Average Daily ON Time Data Sum =

$$\frac{1}{\#tstats} \times \frac{1}{30}\left(\sum_{tstat=1}^{\#tstats} \text{ON Time Data Sum}(tstat)\right)[\text{minutes}] \quad (3.44)$$

The average ON Time error is the average of all the ON time errors individually calculated for each model:

$$\text{Average ON Time Error} = \frac{1}{\#tstats}\left(\sum_{tstat=1}^{\#tstats} \text{ON Time Error}(tstat)\right)[\text{minutes}] \quad (3.45)$$

Thermostat data was collected from 300 anonymous homes in 83 zip codes and five of the seven ASHRAE climate zones. Houses were disqualified if they were missing too much data, appeared to have a major disconnect between the target setpoint temperature and the controlling of the unit, or reversed heating and cooling mode too many times in the sampling period. The first batch of test data used was reduced from 300 to 154 thermostats from May 15 to June 15, 2012. The majority of the thermostats in this set were operating in cooling mode, but there were several heating or had the HVAC turned OFF. The second test set was comprised of 210 thermostats using July 1 to July 31, 2012 data, and these were all in cooling mode. The third was a set of 182 thermostats from January 1 to January 31, 2013 all in heating mode. Figure 3.7 shows the location of the cities where the thermostat data was collected. The figure also displays the seven relevant ASHRAE climate zones.
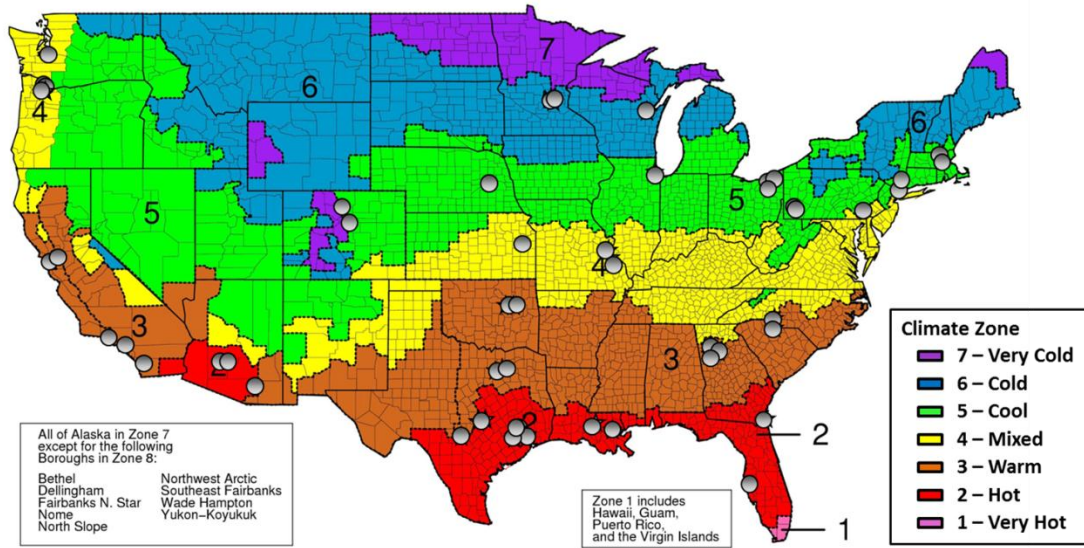
Figure 3.7: Modified ASHRAE climate zone map with the locations of the sampled thermostats (ANSI/ASHRAE 2006).

Houses were then further subdivided by the resolution of weather data used in the model (5-minute or 1-hour). Additionally, the average degrees of setpoint changes (ADSPC) in a day was used to classify the system as having static or dynamic setpoints. This is calculated by summing up all the setpoint changes over 0.56 °C (1 °F) in the 30 day set:

$$ADSPC = \frac{1}{30} \sum_{day=0}^{30} \sum_{t=1}^{288} \left| T_{\text{setpoint}}^{t} - T_{\text{setpoint}}^{t-1} \right| [°C] \tag{3.46}$$

The performance metrics in each category are compared to other relevant categories to determine if differences are statistically significant. The t-test is employed to test two samples and Analysis of Variance (ANOVA) is performed if more than two groups are being compared (Walpole 2007). Statistically significant differences require a p-value lower than 0.05.

3.2.4 Model Performance and Discussion

The following section is divided into two subsections: examples of individual models and analysis on the aggregation of models. The average time to generate a coefficient model and compute all the performance metrics for a house on a dual core 2.33 GHz PC with 8 GB of RAM was less than 90 seconds. Five to ten seconds of this could be attributed to returning data from Amazon's Dynamo Database (Amazon 2013). In a GA run and analysis function the model function was called 25060 times (100 generations, 25 in a population, 10 spider days evaluated; 30 days with model and 30 days control mode) so the average time to compute a single day model was 3 milliseconds. Assuming data from the thermostat does not need to be saved locally (there exists a method to return data from the thermostat when called), the data required to characterize a home can be as small as 200 bytes (15 decimal double values, ID string, and zipcode).

The following examples are used to demonstrate how the Thermostat Building Energy Models compare to the data that trained them. In the plots the solid black lines represent the indoor air temperature data measured by the thermostats with the black shading identifying when the HVAC system was ON. The orange dashed lines show the indoor temperature that the GA determined best matched the data. The red and blue solid lines are the thermostat setpoints for cooling and heating mode respectively. The green dotted lines and shading shows the modeled indoor temperature and predicted ON time from the control simulations. The coefficients for that particular house model are also displayed. Figure 3.8 shows two sample house models on days from the May-June set. Figure 3.8a depicts a house in Charlotte, NC where the model matches the indoor temperature to within 0.44 °C and predicts it will

have the AC ON for 520 minutes compared to the actual 490 minutes on this day.
Figure 3.8b shows a house in Portland, OR that matches the indoor temperature to
within 0.35 °C and predicts the heating will be ON for 890 minutes for this day
compared to the actual 825 minutes. Note the irregular setpoint spikes at 17:00 the
model has to tolerate.
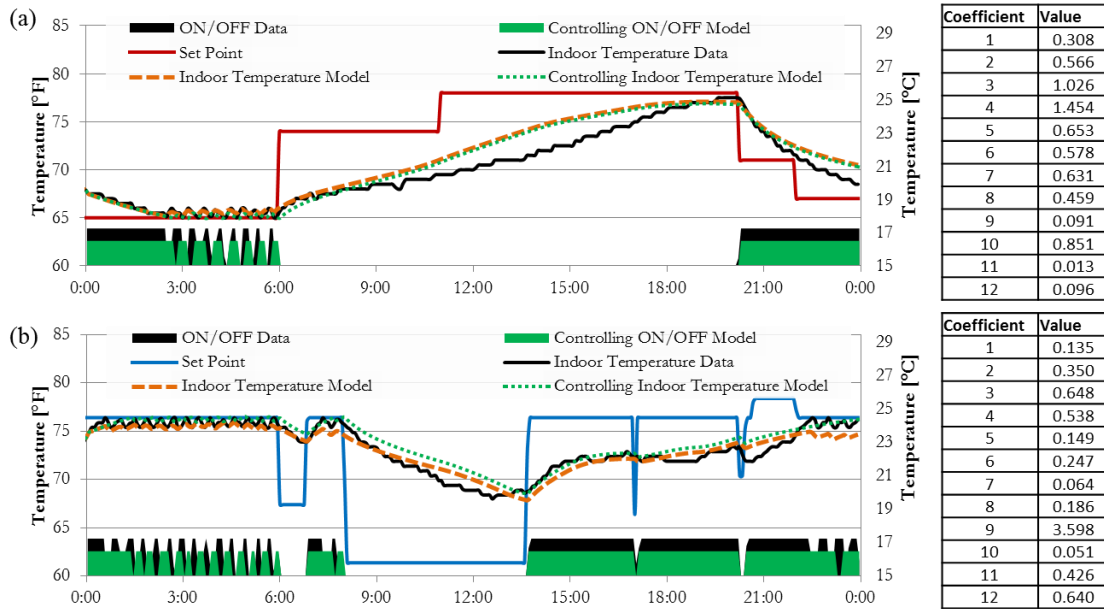
Figure 3.8: (a) Model performance for the indoor temperature and AC ON time on
June 10, 2012 for a house in Charlotte, NC. (b) Model performance for the indoor
temperature and heating ON time on May 22, 2012 for a house near Portland, OR.

Figure 3.9 shows two sample house models on days from the July dataset. Figure
3.9a depicts a house in Tucson, AZ where the model matches the indoor temperature
to within 0.45 °C and predicts it will have the AC ON for 380 minutes compared to
the actual 412 minutes on this day. Figure 3.9b shows a house in New Orleans, LA
that matches the indoor temperature to within 0.25 °C and predicts the cooling will be
ON for 555 minutes for this day compared to the actual 619 minutes. This house is

shown as an example with a flat daily setpoint; the ADSPC is equal to 0 °C.
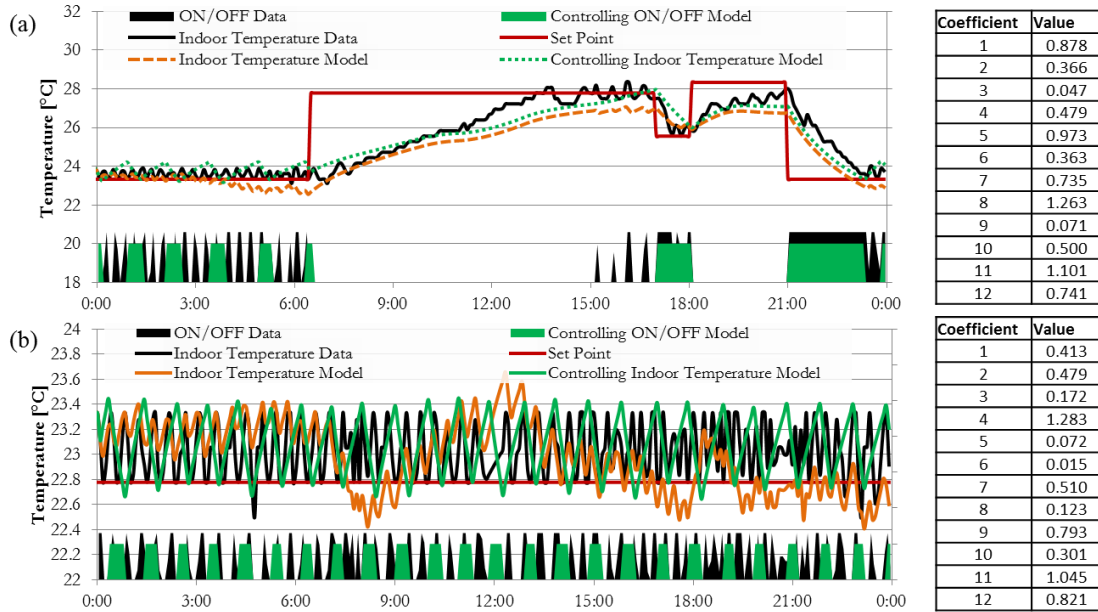
Figure 3.9: (a) Model performance for the indoor temperature and AC ON time on July 22, 2012 for a house in Tucson, AZ. (b) Model performance for the indoor temperature and heating ON time on July 8, 2012 for a house near New Orleans, LA.

Figure 3.10 shows two example house models on days from the January dataset. Figure 3.10a depicts a house in Lowell, MA where the model matches the indoor temperature to within 0.56 °C and predicts it will have the heat ON for 438 minutes compared to the actual 455 minutes on this day. Figure 3.10b shows a house in St. Louis, MO that matches the indoor temperature to within 0.41 °C and predicts the heating will be ON for 210 minutes for this day compared to the actual 199 minutes. This model is shown to illustrate how despite a quality overall model, modeling the hysteresis incorrectly can result in failing to predict HVAC system cycles (6:00-9:00).

53

Coefficient	Value
1	1.725
2	0.021
3	0.001
4	2.223
5	0.163
6	0.140
7	0.112
8	0.026
9	4.510
10	0.075
11	0.449
12	0.022

Coefficient	Value
1	1.676
2	0.017
3	0.002
4	0.842
5	0.144
6	0.176
7	0.082
8	0.027
9	1.726
10	0.300
11	0.210
12	0.230
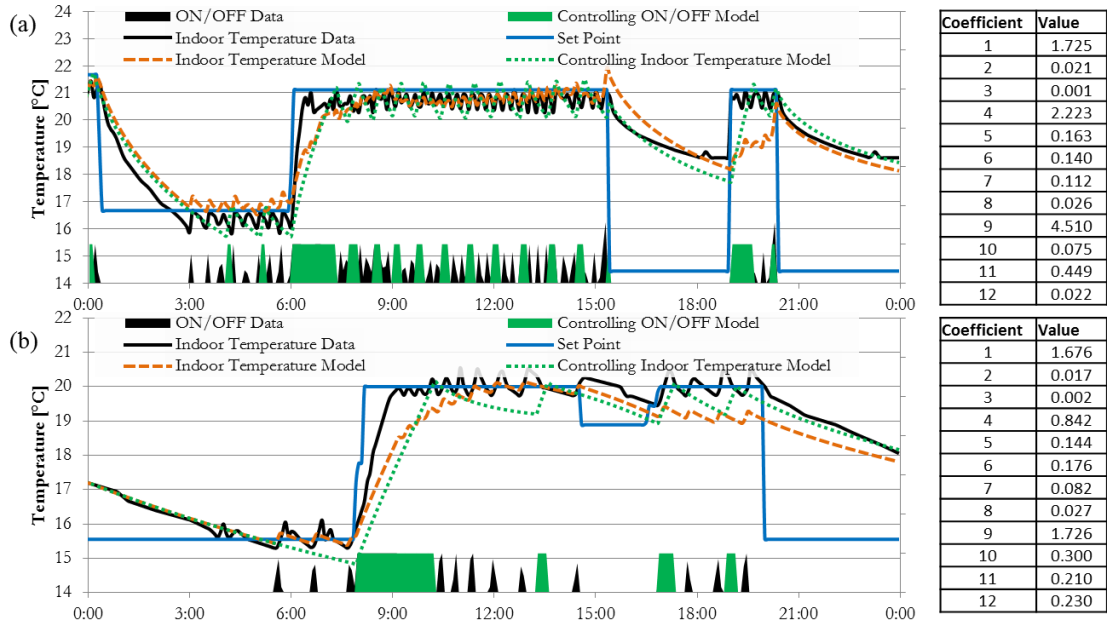
Figure 3.10: (a) Model performance for the indoor temperature and heating ON time for a house in Lowell, MA on January 15, 2013. (b) Model performance for the indoor temperature and heating ON time for a house in St. Louis, MO on January 22, 2013.

Viewing the model on the ten spider days is another way to visualize the performance of the models. Figure 3.11 plots the ten spider days for the May-June dataset for a house in Houston, TX. The model matches the indoor temperature to within an average of 0.45 °C and predicts the air conditioning will be ON for 1330 minutes for this 10-day set compared to the actual 1210 minutes.
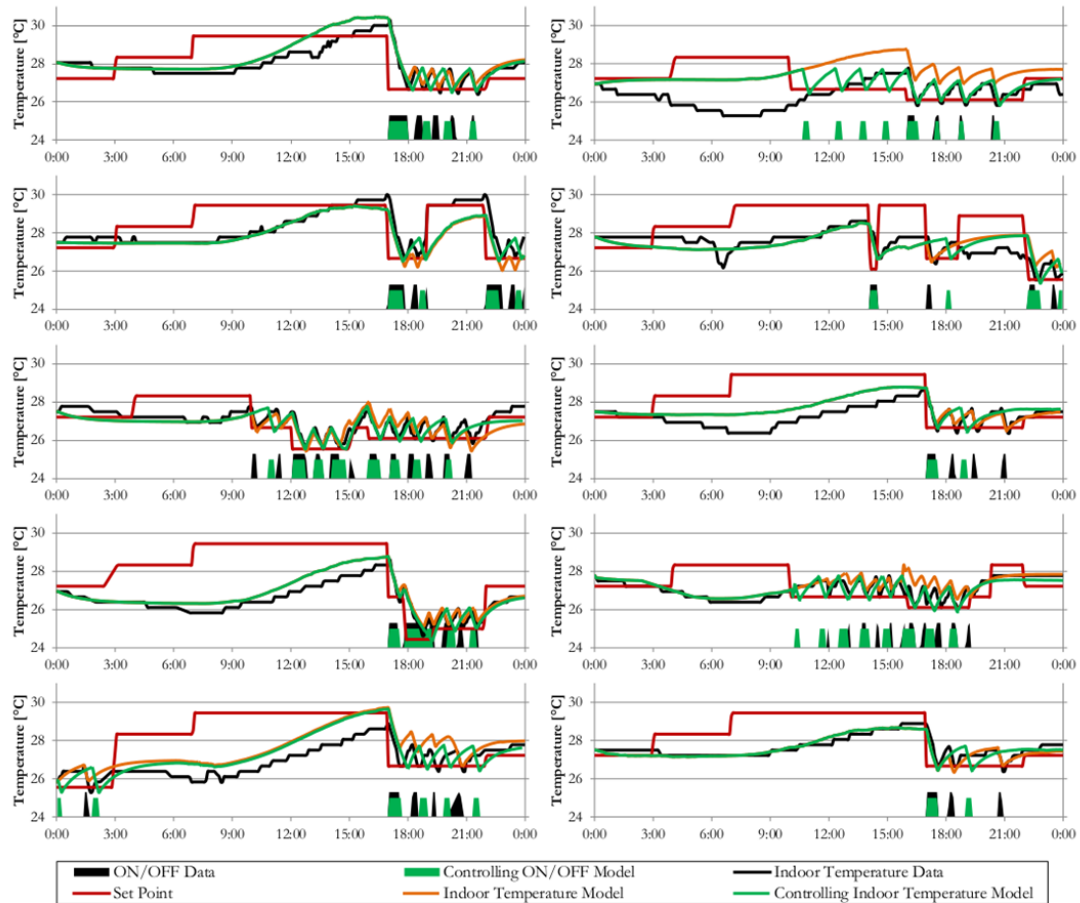
Coefficient	1	2	3	4	5	6	7	8	9	10	11	12
Value	1	0.0261	1.3133	0.0893	0.1752	0.0355	0.0139	0.6372	1.3877	0.0187	0.0818	0.2512

Figure 3.11: Model performance for the indoor temperature and predicted AC ON time on the 10 spider days from May 15, 2012 to June 15, 2012 for a house in Houston, TX.

Figure 3.12 plots the ten spider days for the July dataset for a house in Minneapolis, MN. The model matches the indoor temperature to within an average of 0.42 °C and predicts the air conditioning will be ON for 2745 minutes for this 10-day set compared to the actual 2812 minutes.

55
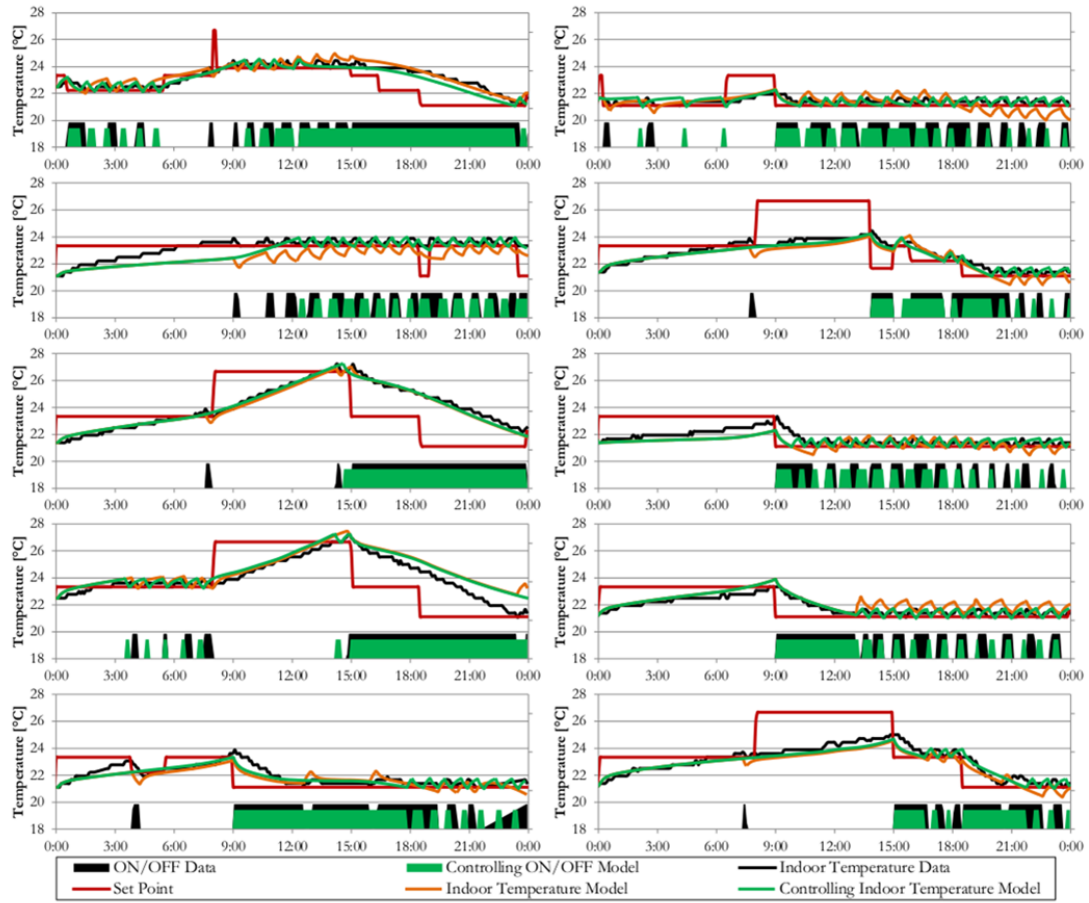
Coefficient	1	2	3	4	5	6	7	8	9	10	11	12
Value	1.055	0.544	0.359	0.531	0.821	0.339	0.091	0.184	0.615	0.023	0.100	0.093

Figure 3.12: Model performance for the indoor temperature and predicted AC ON time on the 10 spider days from July 1, 2012 to July 31, 2012 for a house in Minneapolis, MN.

Figure 3.13 plots the ten spider days for the January dataset for a house in Tacoma, WA. The model matches the indoor temperature to within an average of 0.44 °C and predicts the air conditioning will be ON for 5770 minutes for this 10-day set compared to the actual 5912 minutes. The majority of the days in this set had static setpoints, but the model was still able to match the temperature when the setpoints changed. This plot also shows how an average temperature error of 0.44 °C

56

can be observed in a set with individual errors as large as 1.7 °C.

Coefficient	1	2	3	4	5	6	7	8	9	10	11	12
Value	1.771	0.878	0.690	0.797	0.523	0.331	0.119	4.077	0.035	0.357	0.018	6.075
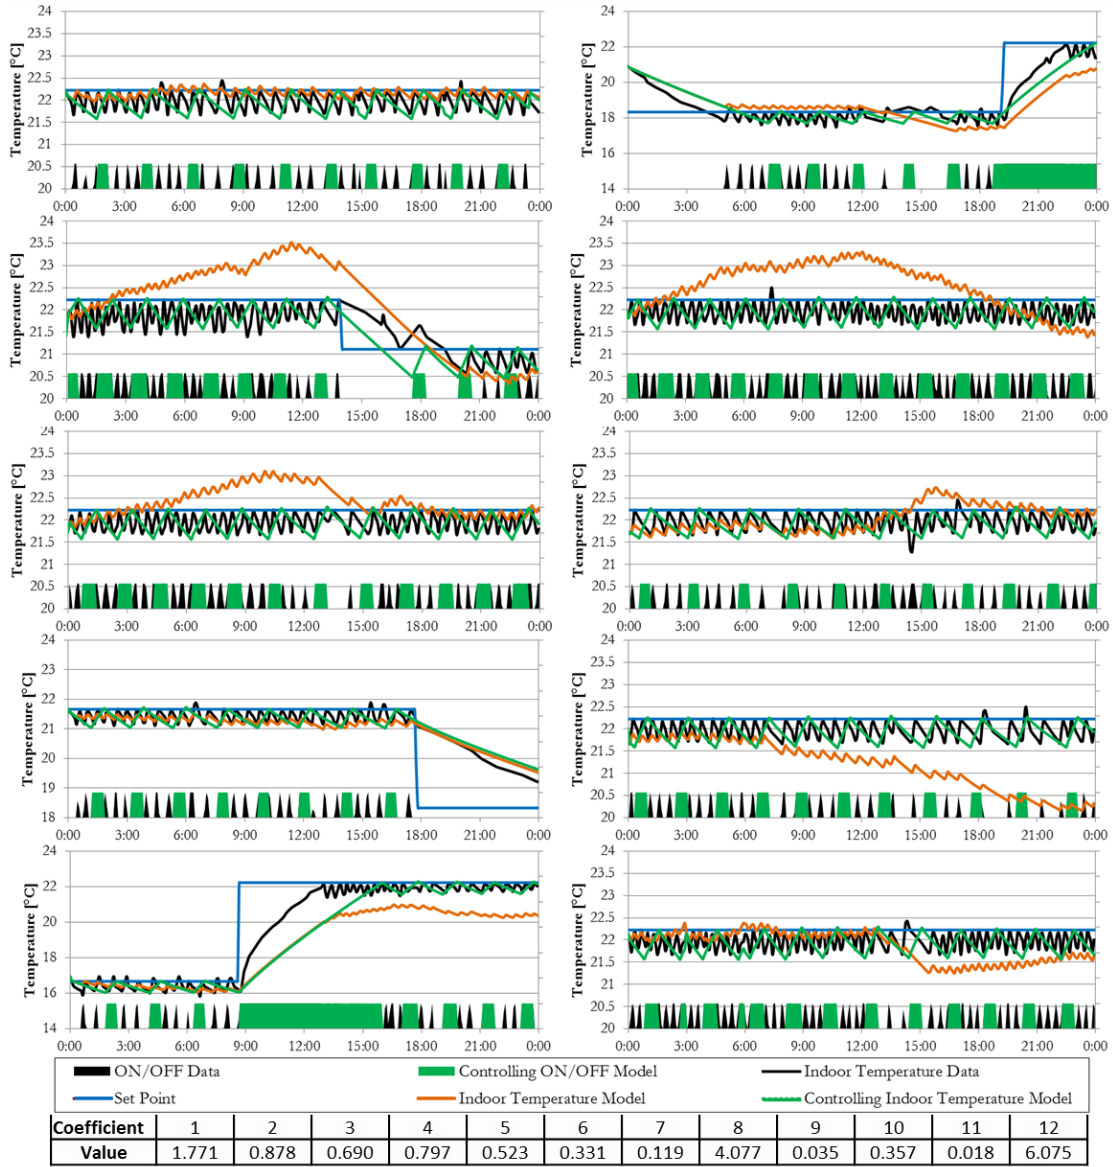
Figure 3.13: Model performance for the indoor temperature and predicted heating ON time on the 10 spider days from January 1, 2013 to January 31, 2013 for a house in Tacoma, WA.

The averaged RMS error for the indoor temperature over the entire May-June set was 0.44 °C. This is an excellent agreement given the 0.28 °C temperature measurement resolution. The model predicts 320 minutes of daily ON time for the

average house where the actual is only 310 minutes, an ON time error of 3.1%. Figure 3.14 illustrates how the majority of houses predicted the ON time sum within 20% of the actual, but the average is still an under-prediction.
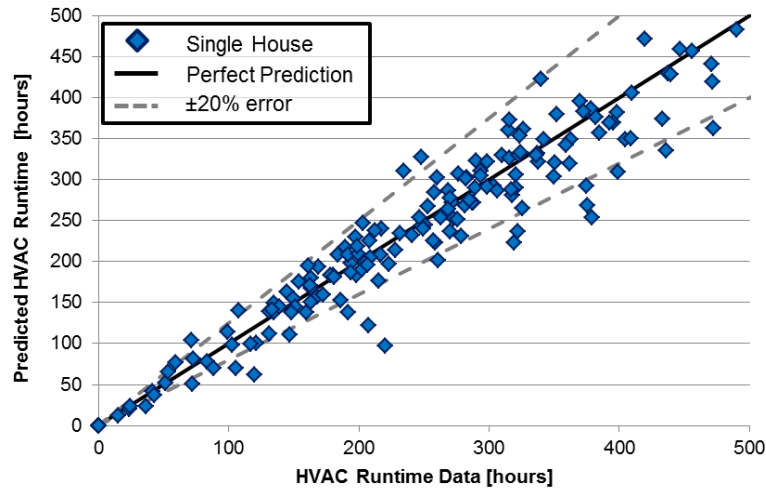
Figure 3.14: Monthly predicted ON time sum verses the ON time data sum for all the houses in the May-June dataset.

While small, there are several likely sources of this systematic error including the time resolution in the ON time data and the thermodynamic model itself. It is even worth considering that the rudimentary filters used to remove uncontrollable days were not sufficient. For most of the country late May to early June is a shoulder season and a popular time to open windows. This was observed in numerous houses yet it is not predictable by the model. Figure 3.15 plots the average daily HVAC duty cycle observed in the data and predicted for the entire set of models from May 15 to June 15, 2012. The shading is used to show when the data and model disagree. Blue shading reflects the models under predicting the duty cycle average and the red is when the model over predicted.
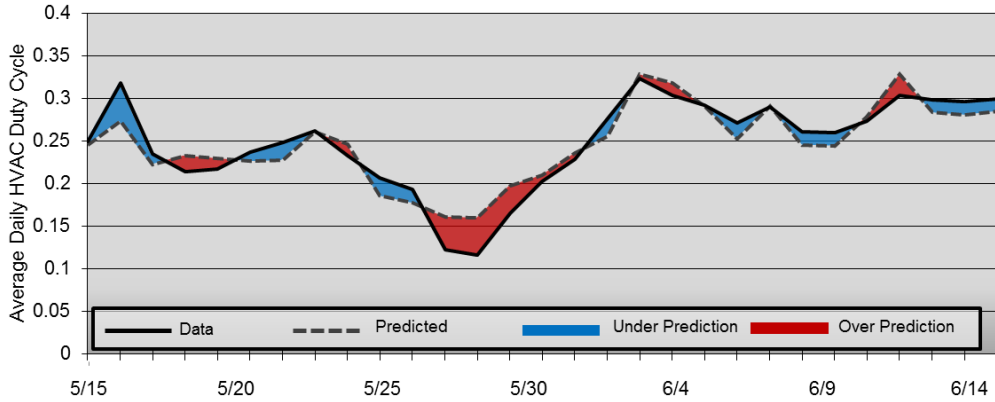
Figure 3.15: Average daily HVAC duty cycle for the May-June dataset.

The all-air-conditioning July dataset averaged a RMS temperature error of 0.43 °C. Over the entire set the models netted an average daily runtime of 510 minutes compared to the actual 502 minutes. This was only a 1.2% over-prediction in runtime. Figure 3.16 shows how only a few houses averaged a prediction in the ON time sum worse than 20%.
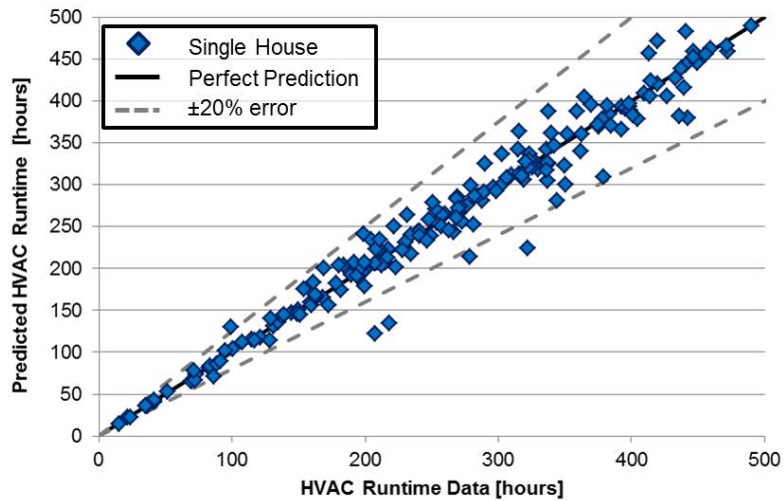
Figure 3.16: Monthly predicted ON time sum verses the ON time data sum for all the houses in the July dataset.

Figure 3.15 plots the average daily HVAC duty cycle observed in the data and

59

predicted for the entire set of models from July 1 to July 31, 2012.
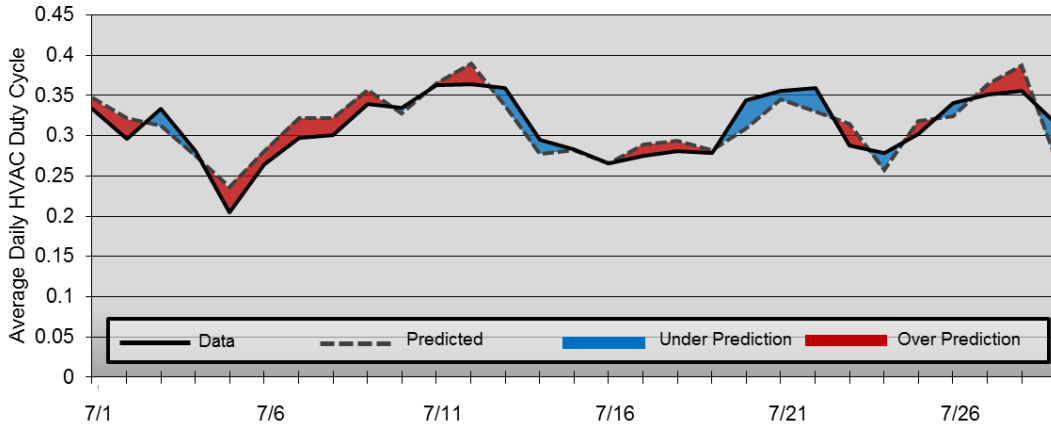
Figure 3.17: Average daily HVAC duty cycle for the July dataset.

The January heating dataset averaged a RMS temperature error of 0.47 °C. Over the entire set the models netted an average daily runtime of 249 minutes compared to the actual 250 minutes. This was only a 0.5% over-prediction in runtime. Figure 3.16 shows how only a few houses averaged a prediction in the ON time sum worse than 20%.
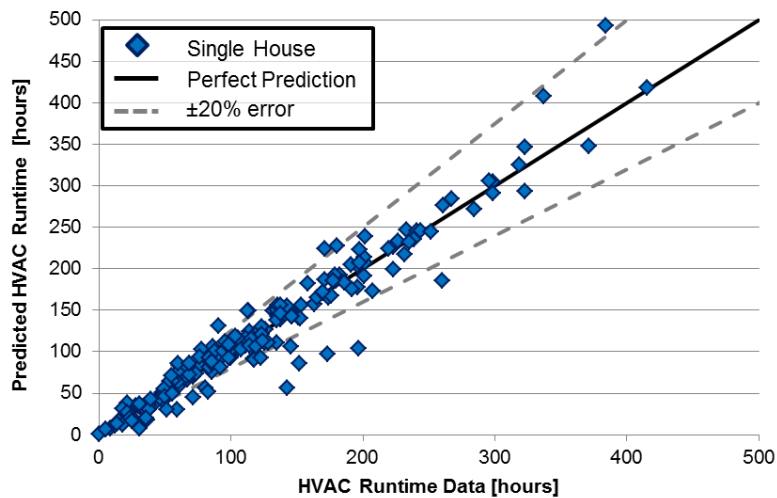
Figure 3.18: Monthly predicted ON time sum verses the ON time data sum for all the houses in the January dataset.

Figure 3.19 plots the average daily HVAC duty cycle observed in the data and predicted for the entire set of models from January 1 to January 31, 2013.
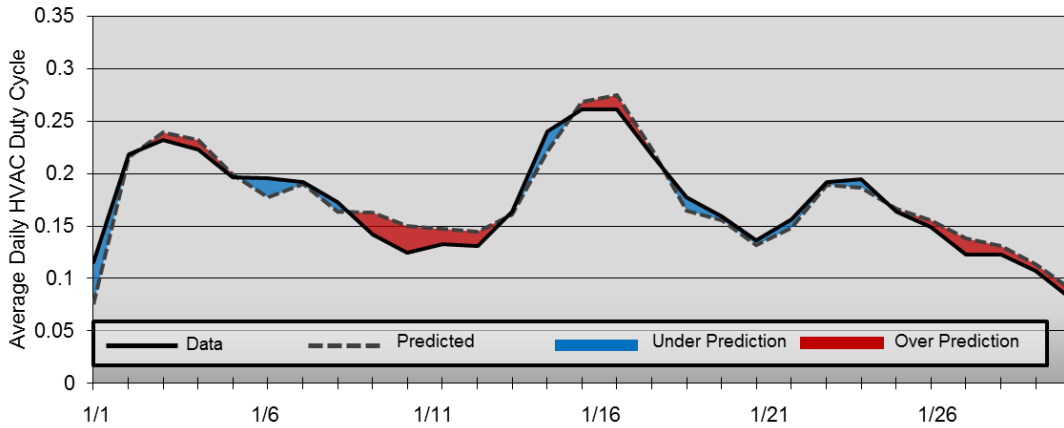
Figure 3.19: Average daily HVAC duty cycle for the July dataset.

The performance of the different model classifications are summarized in the following section. Table 3.4 summarizes the three seasonal tests.

Table 3.4: Summary of the model performance of the three datasets

Classification	Number In Sample	Average Daily Model Predicted ON Time Sum[mins]	Average Daily Actual ON Time Data Sum[mins]	Average ON Time Error [%]	Indoor Temperature RMS [°C]
May-June 2012	154	320	310	2.7	0.44 ± 0.18
July 2012	223	510	502	0.6	0.43 ± 0.21
January 2013	182	249	250	0.1	0.47 ± 0.16

Both heating and cooling models were tested in the three datasets. The single-factor ANOVA test produced a p-value of 0.28 when comparing the ON time error between the different sets of models. This is larger than the critical value of 0.025 and determines that there is not a difference between any of the seasonal models.

The May-June data set was tested training models using both the full 30-day set and only the 10 extreme spider day set. Table 3.5 lists the performance metrics and Figure 3.20 plots the predicted ON time sum verses the actual data of each model

type.

Table 3.5: Models generated with the full month and only spider day training sets for the May-June dataset

Classification	Number In Sample	Average Daily Model Predicted ON Time Sum[mins]	Average Daily Actual ON Time Data Sum[mins]	Average ON Time Error [%]	Indoor Temperature RMS [°C]
Full month	154	318	310	3.7	0.41 ± 0.22
Spider Set	154	320	310	2.7	0.44 ± 0.18
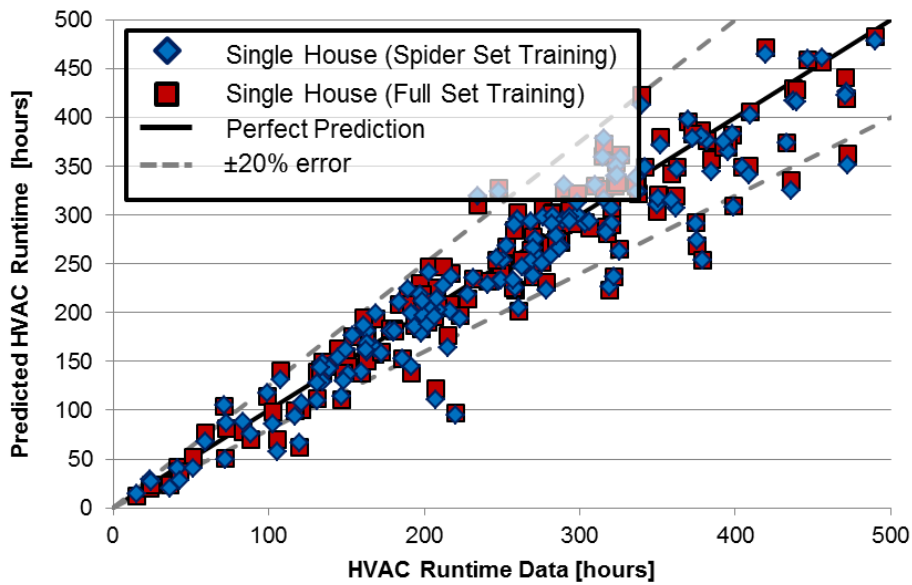
Figure 3.20: Monthly predicted ON time sum verses the ON time data sum for all the houses in the May-June dataset trained using both the full 30 day or spider day set.

Only a marginal improvement in the indoor temp RMS was observed when the models were trained with the entire 30 day set instead of the 10 spider days. The paired t-test p-value was calculated to be 0.33, concluding there is not a significant difference between the methods. Therefore training using only the spider method is preferred because it reduces the computational effort in generating models.

The three datasets are combined for the following analyses to test if there is a

difference between climate zone, weather observation time resolution, and thermostat setpoint variety. The performance metrics from the models in each of the 5 climate zones tested is shown in Table 3.6.

Table 3.6: Performance of the models classified by ASHRAE climate zone

Classification	Number In Sample	Average Daily Model Predicted ON Time Sum[mins]	Average Daily Actual ON Time Data Sum[mins]	Average ON Time Error [%]	Indoor Temperature RMS [°C]
Climate Zone 2	191	421	436	2.2	0.42 ± 0.16
Climate Zone 3	188	397	388	-0.1	0.44 ± 0.19
Climate Zone 4	69	412	410	-0.1	0.45 ± 0.18
Climate Zone 5	81	413	431	3.7	0.42 ± 0.18
Climate Zone 6	22	546	521	-1.4	0.39 ± 0.12

Climate Zone 2 produced the largest collection of thermostats models, and the least were modeled from Climate Zone 6. The single-factor ANOVA test produced a p-value of 0.27 when comparing the ON time error between the different sets of models. Therefore, there is not an observed difference in the ability to produce quality models from any climate zone.

An important test for the weather provider is to determine if hourly weather data is sufficient in generating quality models. The performance metrics between models run using hourly and 5-minute weather observations are shown in Table 3.7.

Table 3.7: Performance of the models classified by weather data observation time resolution

Classification	Number In Sample	Average Daily Model Predicted ON Time Sum[mins]	Average Daily Actual ON Time Data Sum[mins]	Average ON Time Error [%]	Indoor Temperature RMS [°C]
Hourly weather observations	383	421	429	1.7	0.41 ± 0.14
5-minute weather observations	168	407	410	0.8	0.43 ± 0.18

There were more than twice the number of stations trained with hourly weather observation data than 5-minute, but both averaged an ON time error less than 2%. The unpaired t-test p-value was calculated to be 0.51. This suggests that there is no difference between training models with either weather observation time resolution.

The ADSPC metric was designed to separate user thermostat schedules that did not change setpoints throughout the day with ones that did. An example of two extremes would be Figure 3.9b (static ADSPC = 0 °C) and Figure 3.10a (dynamic ADSPC = 26.5 °C). The average ADSPC for the entire set of thermostats was 5.33 °C with a median of 4.0°C and standard deviation of 5.21 °C. Models are classified as having static setpoints if their ADSPC is lower than half the standard deviation below the average (3.72 °C). Dynamic setpoints are ones above half the standard deviation above the average (7.93 °C). Table 3.8 shows the performance metrics of these two classifications.

Table 3.8: Performance of the models classified by setpoint dynamics

Classification	Number In Sample	Average Daily Model Predicted ON Time Sum[mins]	Average Daily Actual ON Time Data Sum[mins]	Average ON Time Error [%]	Indoor Temperature RMS [°C]
Static setpoints	220	464	480	3.5	0.34 ± 0.15
Dynamic setpoints	138	355	349	-2.7	0.54 ± 0.18

More stations were classified as having static setpoints than ones with dynamic setpoints. The average ON time error for the static setpoints was 6.2% higher than the dynamic models. The unpaired t-test p-value was calculated to be 0.0003, which is below the 0.025 critical value needed to reject the null hypothesis.

Therefore, this test concludes that there is a difference in how a model will perform based on the dynamics of the setpoint schedule in the training data. Models from static setpoint schedules tend to under predict the ON time and models from dynamic schedules over predict. It is speculated that this due to the internal thermal mass being too large of a factor. The average duty cycles in Figure 6.6, Figure 6.8, and Figure 6.10 all show that the average system continues being ON more after a peak in the actual data starts to reverse.

The sensitivity of the performance results from the GA was tested using a subset of 10 random houses from the July 2012 dataset. The building model coefficients for each house were determined 10 different times using the C# GA function to see how much the model performance would change. The average, standard deviation, and minimum indoor temperature RMS error of the 10 runs per house are given in Table 3.9. The coefficient set that produced the minimum error was then further evaluated at coefficient values 0.01 higher and lower. This was to simulate the gradient at the GA solution. Table 3.9 also contains the lowest error found by this search.

Table 3.9: Sensitivity of the GA fitness

	Indoor Temperature RMS Error [°C]			
House #	Average	STDEV	Minimum	Gradient Search
1	0.450	0.016	0.432	0.431
2	0.398	0.004	0.393	0.390
3	0.891	0.004	0.881	0.878
4	0.377	0.002	0.375	0.369
5	0.423	0.002	0.420	0.419
6	0.466	0.002	0.462	0.462
7	0.288	0.003	0.285	0.281
8	0.210	0.001	0.208	0.207
9	0.456	0.015	0.431	0.429
10	0.491	0.005	0.484	0.473
Average	0.445	0.005	0.437	0.434

This analysis shows how the heuristic nature of the GA function produces different results each time the function is used to determine the coefficients. The indoor temperature RMS error does differ, but not by much. The largest standard deviation in the coefficient search was only 0.016 °C or 3.5% of the total error. When searching 10 times the best model was only on average 0.008 °C better than the average. Also further exploring the solution set by evaluating the gradient only improved the error by an average of 0.003 °C. These differences in the coefficient models were insignificant enough to not change the predicted ON time, so the average ON time error remained constant.

These 10 houses were also evaluated using optimization solution functions in MATLAB. Specifically the MATLAB GA, Fminunc (unconstrained gradient based optimization), Fminsearch (non-gradient based optimization search), and Lsqnonlin (non-linear least squares) were used to test how the indoor temperature RMS error would change given a different solution function. The Fminunc and Lsqnonlin functions were given penalties to prevent negative coefficient solutions. Table 3.10 presents the results of this analysis.

Table 3.10: Determining the model coefficients using multiple MATLAB functions

House #	Indoor Temperature RMS Error [°C]			
	GA	Fminunc + Penalty	Fminsearch	Lsqnonlin + Penalty
1	0.445	0.550	0.413	0.412
2	0.398	0.437	0.376	0.370
3	0.887	0.966	0.884	0.880
4	0.372	0.404	0.378	0.378
5	0.418	0.436	0.412	0.408
6	0.461	0.519	0.462	0.458
7	0.281	0.334	0.280	0.277
8	0.202	0.268	0.193	0.190
9	0.452	0.536	0.364	0.361
10	0.486	0.494	0.465	0.465
Average	0.440	0.494	0.423	0.420

This analysis showed that the GA solution technique can be improved. Both the Fminsearch and Lsqnonlin functions produced coefficient models with lower average indoor temperature RMS errors, but it was only a 4.5% improvement.

The results of the gradient search at the end of the GA solution and the improvement using the MATLAB functions indicate that the C# GA function was not obtaining a global or even local optimum. However, the closeness of the indoor temperature RMS errors between these suggests that obtaining the optimum may not be required for accurate predictions. If the thermostat building energy model is to be reproduced, a non-linear least squares solution technique is recommended to determine the coefficients, not GA.

3.2.5 Initial Conclusions

The results from these tests give evidence that most houses can be modeled using the thermostat building energy model successfully. The entire set averaged an error in predicting the ON time of 1.9% and indoor temperature RMS error of 0.44°C. These results are independent of what energy system (AC or heating) was used. The large number of houses and the variety in geographic location and climate in the sample strengthens this claim as well. Separating the models by the ADSPC metric identified that the model can be improved, but the overall predictions were still quite accurate. Further improvement in the accuracy was observed when utilizing coefficient solution techniques other than GA. It is recommended to continue this work in the future with non-linear least squares instead. The price and installation effort may be a barrier to older houses. It is expected that most of the houses tested were newer, but the errors for the true average house should not be significantly different than what was

presented here. Additional plots displaying the duty cycle and model indoor air temperature trends of these datasets can be found in the Appendix.

3.2.6 Energy Consumption Modeling with Thermostats

The previous section explained how internet-connected thermostats could be used to predict when a home's HVAC system would turn ON. To fully predict transient energy consumption, modeling the energy consumption of the HVAC system is required to convert ON time to power/energy. Residential smart power meters are a new resource that is used to solve this problem. This section explains a method to disaggregate the HVAC load from the total load using the internet-connected thermostat data. This method also generates an HVAC power curve that relates the power to outdoor air temperature.

This method requires both smart meter data and an internet-connected thermostat. This is done in an effort to identify spikes in power only associated with the HVAC and not other large appliances. Air conditioners and heat pumps draw power at rates that are dependent on the condensing /evaporating air temperature. See Figure 2.8 for an example of recorded power of an air conditioner verses time; specifically how the load is lower at hour 6 than hour 14 and assuming the morning is cooler outside than the afternoon. Smart meter data is available in the CenterPoint Energy (Houston, TX area) territory in approximately 15-minute intervals. Despite it being named a power meter, the data is reported in aggregate kWh since the last measurement. The thermodynamic house model and internet-connected thermostat readings operate in 5-minute intervals; therefore, these data are aggregated or averaged to match the time resolution of the smart meter data. The same "percentage

of ON time interpolating method" from 2.2.1 is employed to do this. The HVAC runtime signal is filtered to find instances when HVAC goes from being fully OFF for 15 minutes to fully ON for the next 15 minutes or vice versa. This is referred to as a complete ON/OFF switch. When this is identified, the difference in power between the ON and OFF point are recorded and matched with the outdoor air temperature of the ON time step. Figure 3.21 shows there were only three of these observed in the 11-hour data set plotted (the white filled data points signify the observed complete ON/OFF switches).
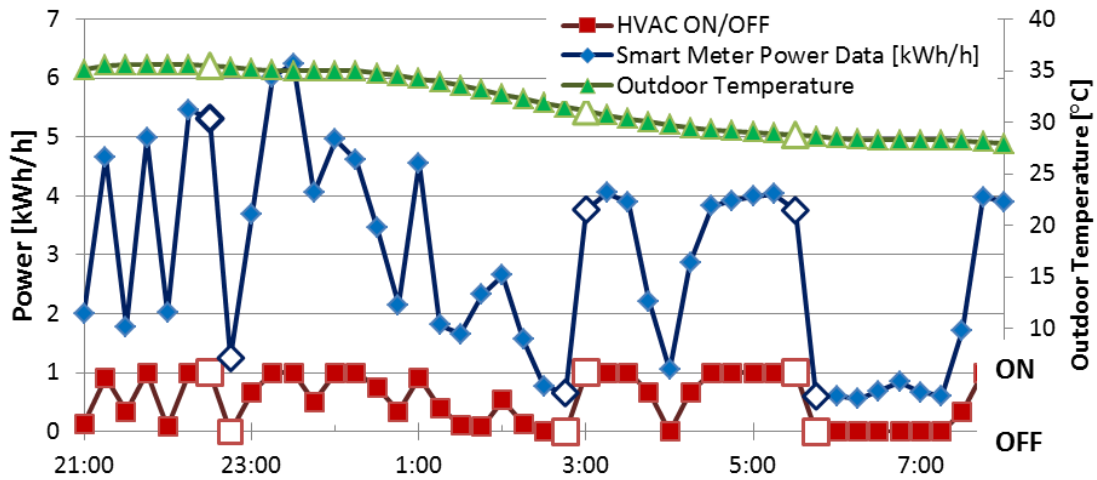
Figure 3.21: Identifying complete ON/OFF switches in the smart meter power data.

These recorded points are assuming that any other power consuming devices in the data are going to be continuing the same pattern through this switch. The recorded points are separated into ten evenly spaced temperature bins by the outdoor temperature in an effort to provide equal weight to the range in temperatures. AC manufactures publish a linear power curve relating the power of the compressor and outdoor air temperature. Therefore, a linear regression of bin averaged power data is employed to develop an approximate power curve for the house's system. Figure 3.22

plots all the observed complete ON/OFF switch power differences and the binned

data along with the linear regression that is used as the power curve for this home's

HVAC system.
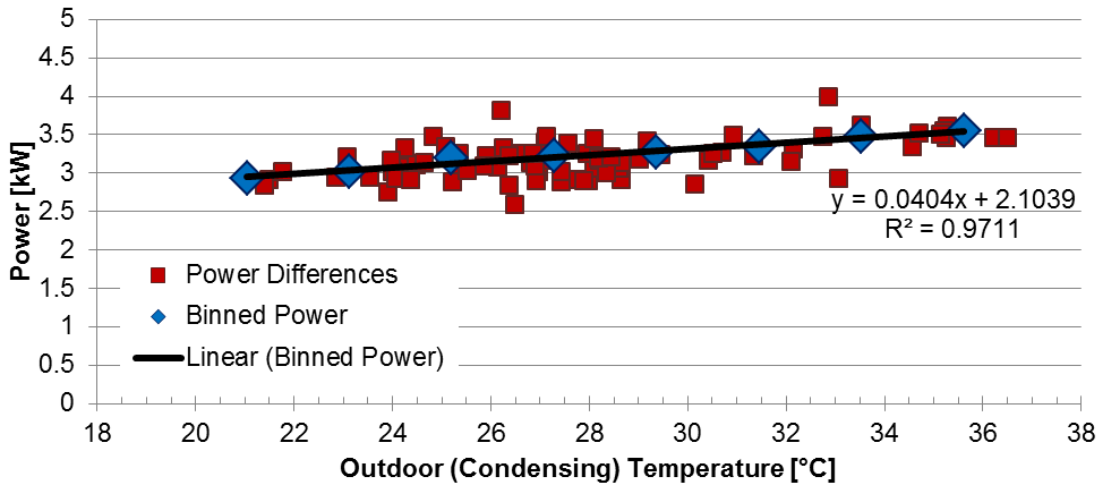
Figure 3.22: Observed and binned power differences used to generate a power curve

for a house in Houston, TX using data from July 2012.

Houses can have multiple thermostats per smart meter signal, and logic was put

in place to eliminate instances where these signals would interfere. In the CenterPoint

dataset there were 144 unique smart meters for 186 thermostats. The average house

was 2729 ft^2, but was reduced to 2118 ft^2 per thermostat by splitting up the area of

multi-zone systems. In ASHRAE climate zone 2 this would require a 4-ton air

conditioner (ANSI/ASHRAE 2006). Six power curves are plotted in Figure 3.23 with

two representing economy and high efficiency 4-ton units (14 and 21 SEER)

manufactured by Carrier, the average over the entire CenterPoint dataset derived from

this method, and two specific examples from houses that have an area to suggest a 4-

ton unit. The one labeled "efficient" matches the 21-SEER Carrier unit quite well and

the one labeled "inefficient" requires an increasingly larger amount of power than the

new 14-SEER unit.
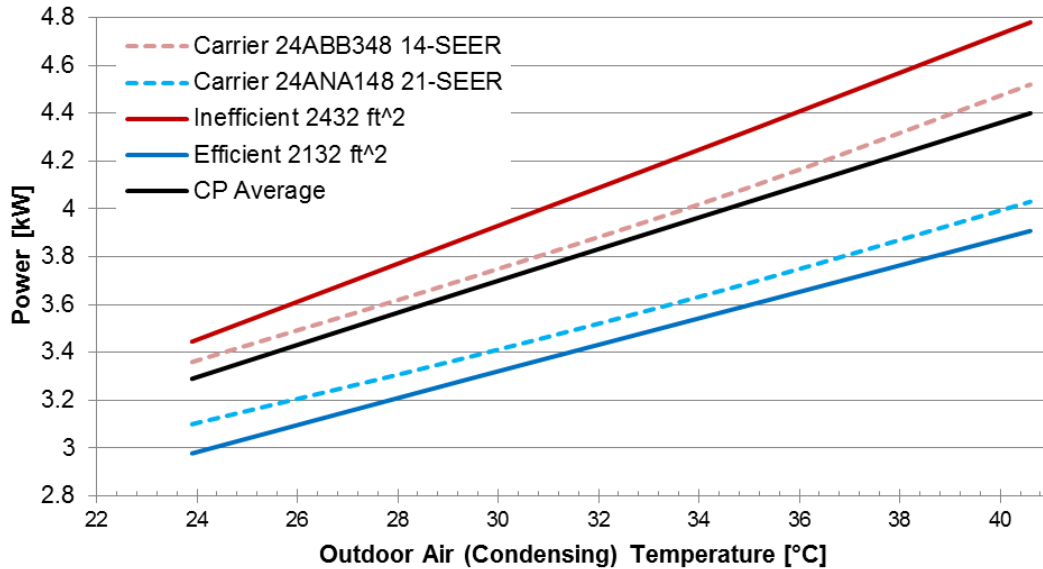
Figure 3.23: Power curves reported by Carrier and observed in the CenterPoint dataset.

The average power curve in the data set was expected to be closer to Carrier's economy unit because most HVAC contractors opt to install the economy unit. It was not expected for the average to be a bit more efficient than the economy unit though. This may be attributed to the suggestion that homeowners who installed these internet-connected thermostats are more concerned with energy savings than the average homeowner, causing them to opt for higher efficiency units.

Once the power curve is identified the power required to run the HVAC can be approximated for past or forecasted future conditions given an outdoor temperature. Figure 3.24 gives an example of how the total and disaggregated HVAC energy data can be visualized for customers.
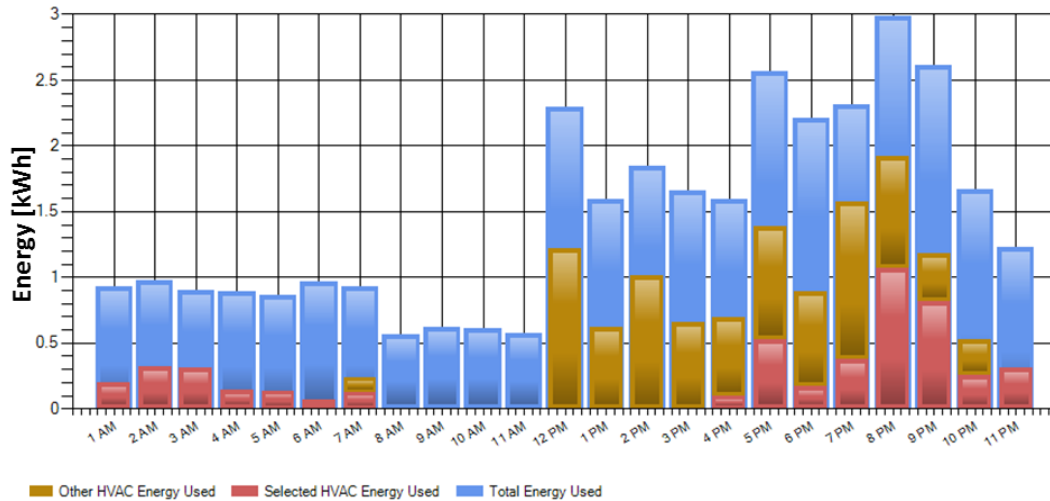
Figure 3.24: The disaggregated energy consumption of two air conditioning units from the total house energy consumption for a house in Houston, TX on September 13, 2012.

Forecasting the entire home energy consumption requires an approximation for the load not associated with the HVAC. This is done by averaging the remaining load after disaggregating the HVAC at every time interval throughout a day. This average load can be added to any forecasted HVAC load to forecast the total load. Further modeling of the non-HVAC load can be performed to improve the final prediction, but averaging was used for this model. Figure 3.25 summarizes the system involved in the energy consumption forecast calculation. Datasets required to train the models are on the left and ones to produce an energy consumption forecast are on the right. Figure 3.26 plots the aggregate daily energy consumption of 250 houses in Houston, TX from July 1 to September 30, 2012. The plot also displays the predicted energy consumption using the thermostat building energy models. This plot is generated using the actual weather observations and Figure 3.27 shows the daily forecasted consumption using forecasted weather data. The forecasted weather data is updated at

0:00, 6:00, 12:00, and 18:00 each day. These plots also use the shading convention of

red for a model over-prediction and blue for a model under-prediction.

Thermostat Building Energy Model

Training Data

Indoor Temperature

Thermostat Data   ON /OFF

Set Point Schedule

Weather Data

12-Coefficent Thermodynamic Building Model

Operation Data

Thermostat Set Point Schedule

Weather Forecasts

HVAC ON (time)

Thermostat ON / OFF Data

Weather Data

Smart Power Meter Data

Power Curve and Average Base Load Model

Weather Forecasts

HVAC Energy (time)

Average Base Load (time)

Energy Consumption Forecast
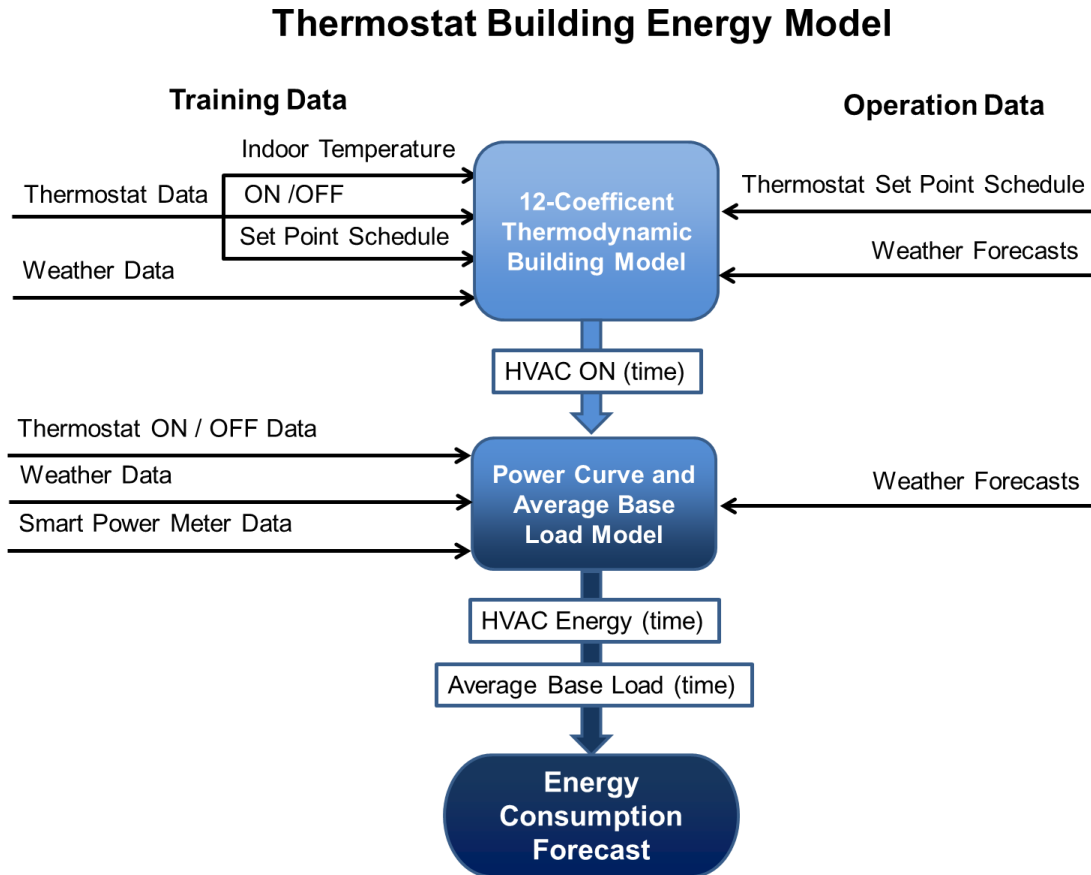
Figure 3.25: System diagram for energy consumption forecasting using the thermostat

building energy model.
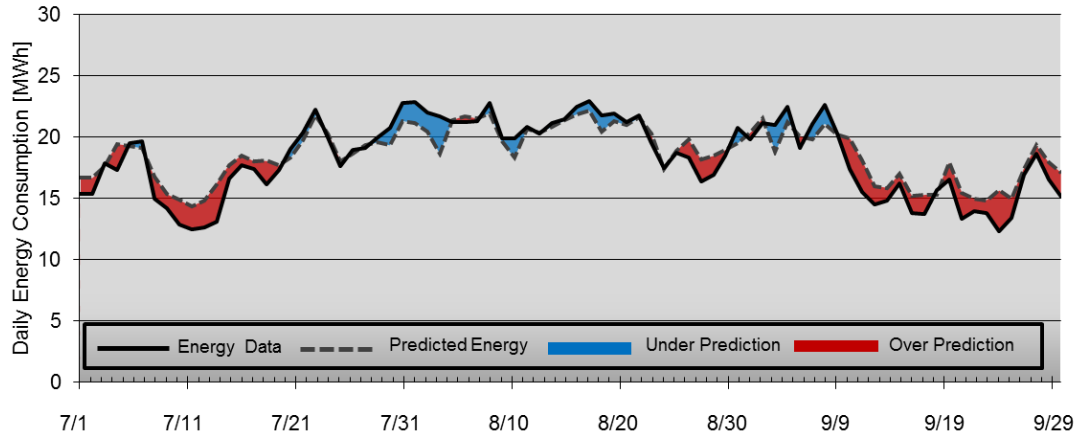
Figure 3.26: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from the summer of 2012 calculated using actual weather observation data.
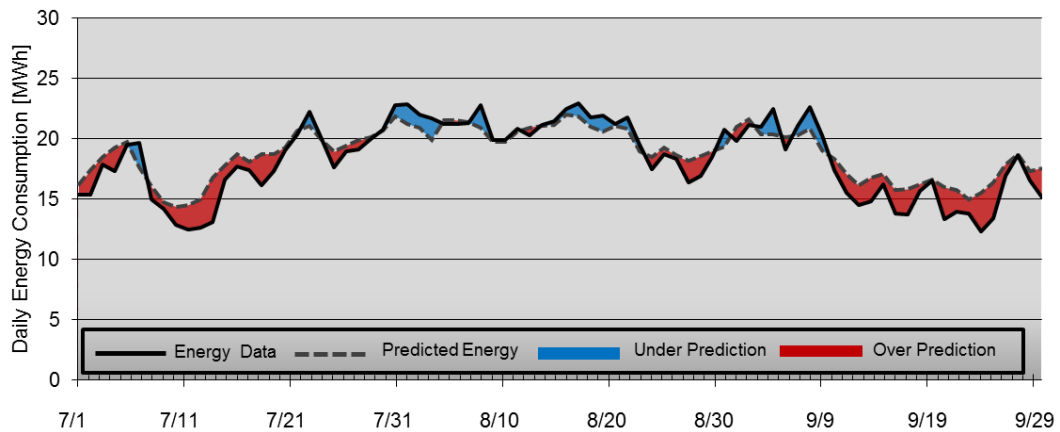
Figure 3.27: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from the summer of 2012 calculated using weather forecast data.

These data are also viewed hourly so see how the forecast compares to the actual consumption throughout the day. Figure 3.28 plots the forecast generated using the actual weather observation data and Figure 3.29 plots the energy forecast generated from forecast weather data. The week of July 18 to July 25, 2012 was chosen to

74

highlight because there was a large range in daily peak consumption.
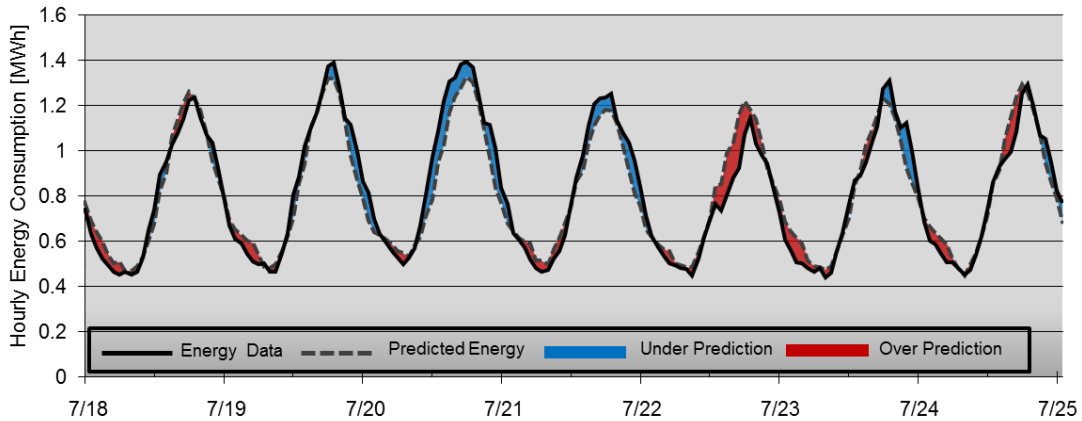
Figure 3.28: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using actual weather observation data.
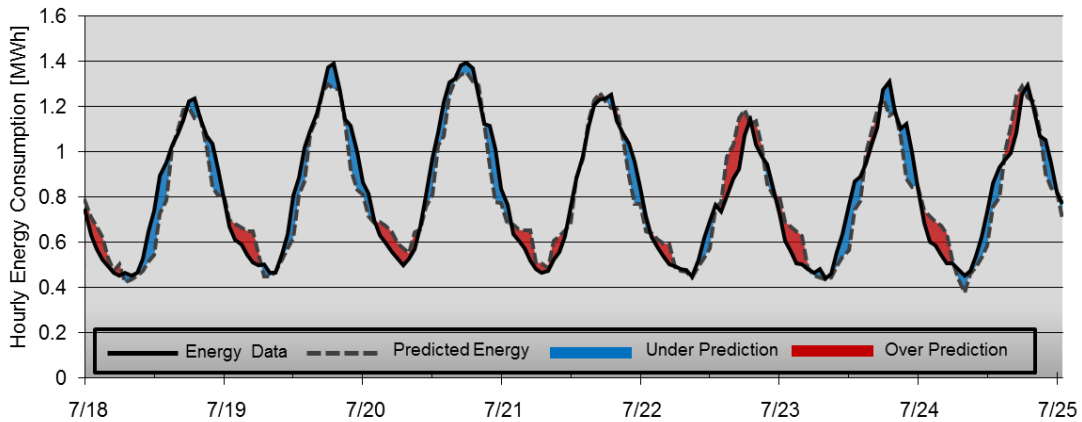
Figure 3.29: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using weather forecast data.

The measured total energy consumed by the 250 houses from July 1 to September 30, 2012 was 922 MWh. The model predicted the consumption would be 930 MWh when using the actual weather observation data and 932 MWh when using

the weather forecasts. The daily RMS error in prediction for the set was 86 MWh for models using weather observations and 89 MWh for models using forecasts. The hourly RMS error for the selected week was 162 MWh when using observations and 191 when using forecasts.

3.3 Smart Meter Only Building Energy Model

3.3.1 Model Derivation

A building energy forecasting model was also developed that operated without the internet-connected thermostat in an effort to see if the forecasting performance could be replicated using only smart meter and weather data. Instead of only correlating the energy consumption to temperature data, like a traditional overall heat transfer method, this method additionally uses the sun loading and wind speed data (ASHRAE 2009). An effective heat transfer from the outdoor to indoor environment metric is derived to do this. The same smart meter data from CenterPoint Energy set in Houston, TX from the summer of 2012 was used to develop this method.

Assumptions were made about the house to reduce the complexity of the thermostat building energy model, but a few new energy sources were introduced as well. The heat flows modeled are conduction through the walls and roof, solar loading on the walls, roof, and through the windows, infiltration, internal heat generation, and energy stored in the thermal mass of the house. Figure 3.30 is a diagram of these energy flows.
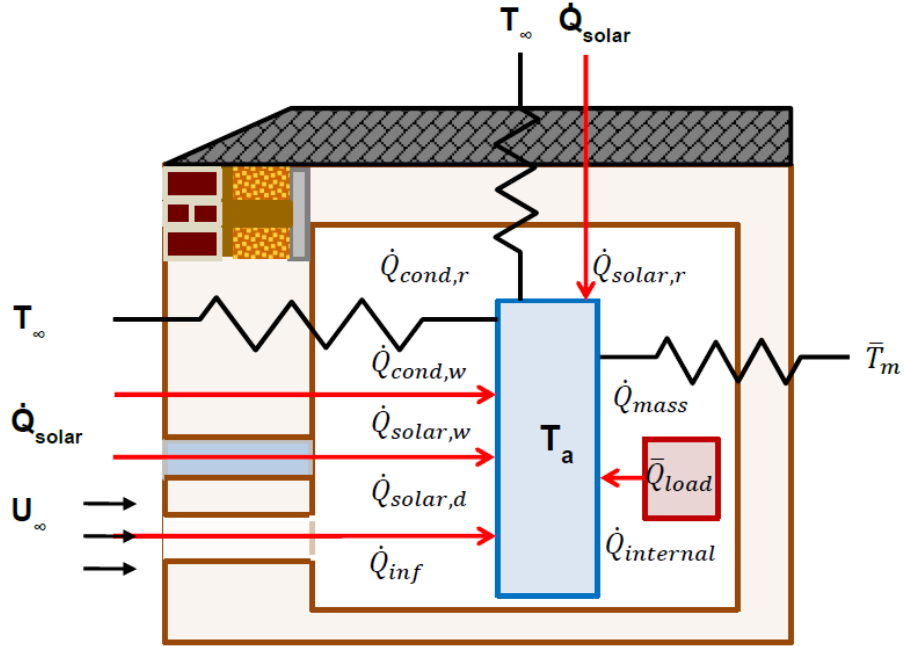
Figure 3.30: Smart meter only building energy model thermoelectric analogy model.

The indoor air temperature was assumed to be a constant 24 °C for every house because it was the most commonly observed setpoint from the summer in the tested dataset. This can be changed to reflect the season or region, and can even be made into a variable. A constant indoor air temperature reduced the need for the heat diffusion differential equation. This was replaced by adding a time constant coefficient (τ) to a simplified conduction equation.

$$\dot{Q}^t{}_{conduction} = \left(A_r \frac{1}{R_r} + A_w \frac{1}{R_w}\right)(T^{t-\tau}_\infty - 24°C) \tag{3.47}$$

It was assumed this time constant was the same for energy flowing through the walls as the roof. The solar terms were modeled using a solar factor (SFW = solar factor walls and SFD = solar factor windows).

$$\dot{Q}^t{}_{solar\ load} = \dot{Q}^{t-\tau}_{solar}\big((A_r + A_w)SFW\big) + \dot{Q}^t_{solar}A_w SFD \tag{3.48}$$

Infiltration was modeled containing both wind driven and general temperature

difference components and an infiltration factor (INF).

$$\dot{Q}^t_{infiltration} = \left(U^t_\infty{}^2 + 1\right)\left(T^t_\infty - 24°C\right)INF \qquad (3.49)$$

Internal heat generation was modeled as a heat factor (IHF) times the average hourly power consumption measured over the entire season.

$$\dot{Q}^t_{internal} = \bar{Q}^t_{load}IHF \qquad (3.50)$$

The thermal mass of the structure was modeled as the difference between the average outdoor air temperature from the last three days and the indoor air temperature times the heat capacity and amount of mass (TMF). This time period was determined through trial and error.

$$\dot{Q}^t_{mass} = (\bar{T} - 24°C)TM \qquad (3.51)$$

The area of the walls and roof are approximated using the floor space of the house. Eq 2.52 gives the overall heat transfer equation with the unknown building parameters and intermediate variables calculated from input data in

Table 3.11.

$$\dot{Q}^t = \left(A_r \frac{1}{R_r} + A_w \frac{1}{R_w}\right)\left(T^{t-\tau}_\infty - 24°C\right) + \dot{Q}^{t-\tau}_{solar}\left((A_r + A_w)SFW\right) +$$

$$\dot{Q}^t_{solar}A_wSFD + \left(U^t_\infty{}^2 + 1\right)\left(T^t_\infty - 24°C\right)INF + \bar{Q}^t_{load}IHF + (\bar{T} - 24°C)TM$$

Table 3.11: Unknown building parameters and intermediate variables in the heat transfer calculation

Building parameter	Variable in Equation	Intermediate Variable	Variable in Equation
Thermal resistance of the roof	R_r	Roof surface area	A_r
Thermal resistance of the walls	R_w	Wall surface area	A_w
Wall solar factor	SFW	Average total load	$\bar{Q}^t_{load}$
Window solar factor	SFD	Average outdoor T	$\bar{T}$
Infiltration factor	INF		
Internal heat factor	IHF		
Thermal Mass	TM		
Diffusion time constant	τ		

For a particular set of building parameters a heat transfer value is calculated for every point in the season. The power measurement at each time step is binned by the heat transfer. These data points are fit to a second-order polynomial using OLS. A second-order polynomial correlation was used because there is a large range in the lower heat transfer bins (-1000 to 2500W) where the average power is relatively

constant. The AC is not expected to cycle as frequently during low heat transfer as it is when conditions heat up (+4000W). When this happens the average load increases. Figure 3.31 plots all the hourly averaged electrical power points versus the heat transfer metric calculated using a set of building parameters. The binned data is used to develop the second-order polynomial correlation between electrical power and heat transfer.
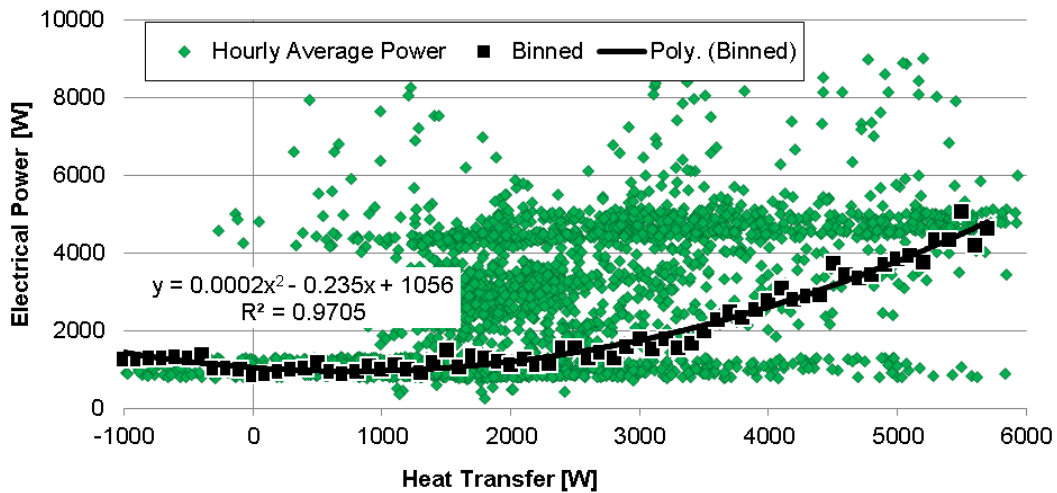
Figure 3.31: Electrical power correlation with the heat transfer metric for a house in Houston, TX.

This correlation is used to predict the electrical energy consumption for every hour in the dataset. The unknown building parameters were determined using the same genetic algorithm function employed to generate the thermostat building energy models. The GA function minimized the sum of the difference between the seasonal energy consumption and that calculated from a heat transfer correlation. Like the thermostat building energy modeling, the GA can be replaced with a more efficient solution technique. The mathematical formulation of the problem is given in Equation 3.53.

$$\min_{c_1, c_2, \ldots c_8} \sum_{day=0}^{92} \sum_{t=0}^{24} \left(P^t - A \left(\dot{Q}^t (c_1, c_2, \ldots, c_8) \right)^2 + B \left(\dot{Q}^t (c_1, c_2, \ldots, c_8) \right) + C \right)$$

$$s.t.:$$

$$c_1, c_2, \ldots, c_8 \geq 0$$

(3.53)

In this formulation the eight building parameters are represented as eight coefficients (c_1, c_2, …, c_8). The 92 days span the entire summer dataset from July 1 to September 30, 2012, and the 24 daily time points are for each hour. The A, B, and C are the coefficients of the heat transfer least squares curve fit correlation, and P^t is the measured electrical power.

The system for training this model and performing the energy consumption forecast calculation is summarized in Figure 3.32.
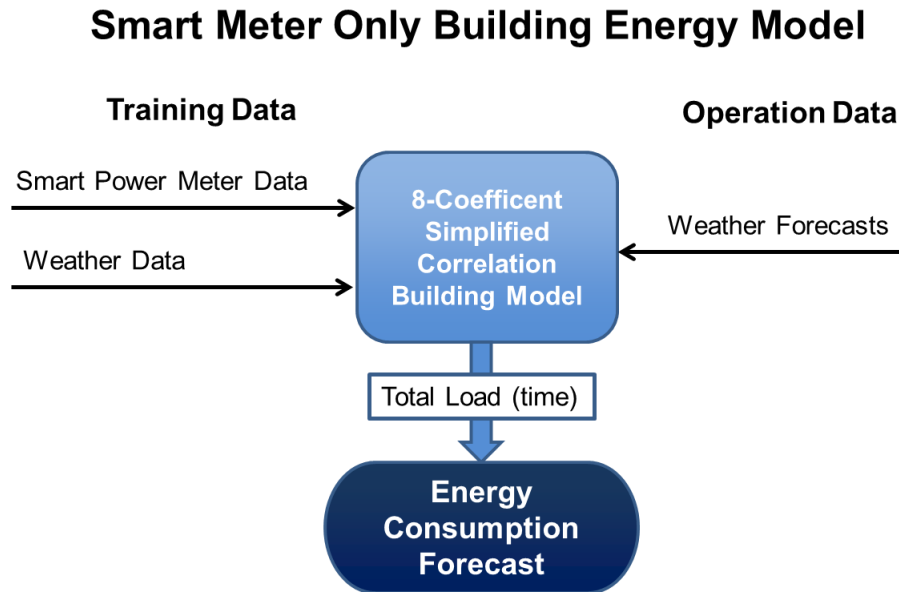
Figure 3.32: System diagram for energy consumption forecasting using the smart meter only building energy model.

3.3.2 Model Performance

Figure 3.33 plots the aggregate daily energy consumption of 250 houses in

81

Houston, TX from July 1 to September 30, 2012. The plot also displays the predicted energy consumption using this system. This plot is generated using the actual weather observations and Figure 3.34 shows the daily forecasted consumption using forecasted weather data. The forecasted weather data is updated at 0:00, 6:00, 12:00, and 18:00 each day.
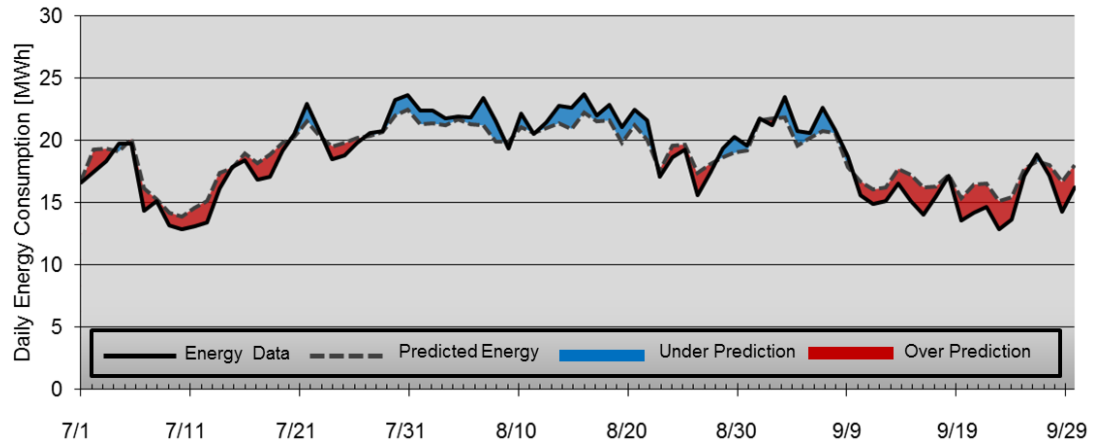
Figure 3.33: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from the summer of 2012 calculated using actual weather observation data.
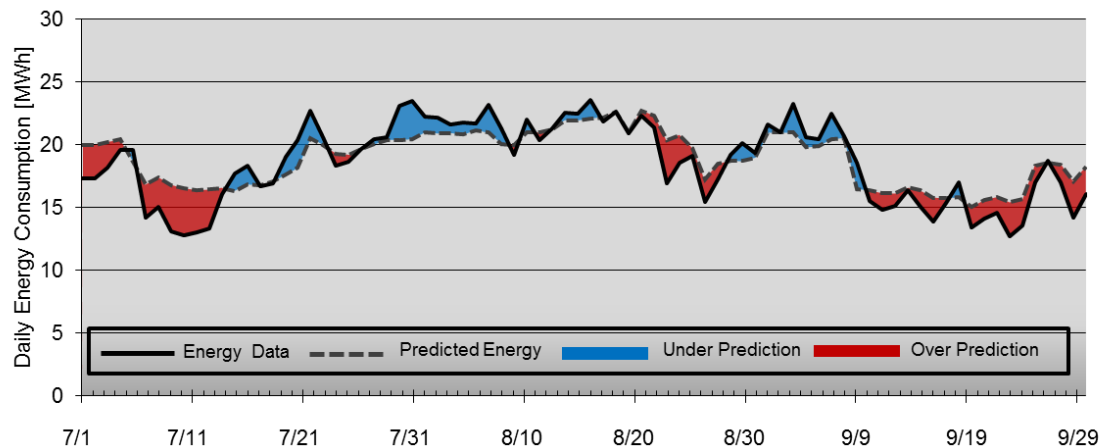
Figure 3.34: Predicted and actual aggregate daily energy consumption for 250 houses in Houston, TX from the summer of 2012 calculated using weather forecast data.

These data are also viewed hourly to see how the forecast compares to the actual consumption throughout the day. Figure 3.35 plots the forecast generated using the actual weather observation data and Figure 3.36 plots the energy forecast generated from forecast weather data.
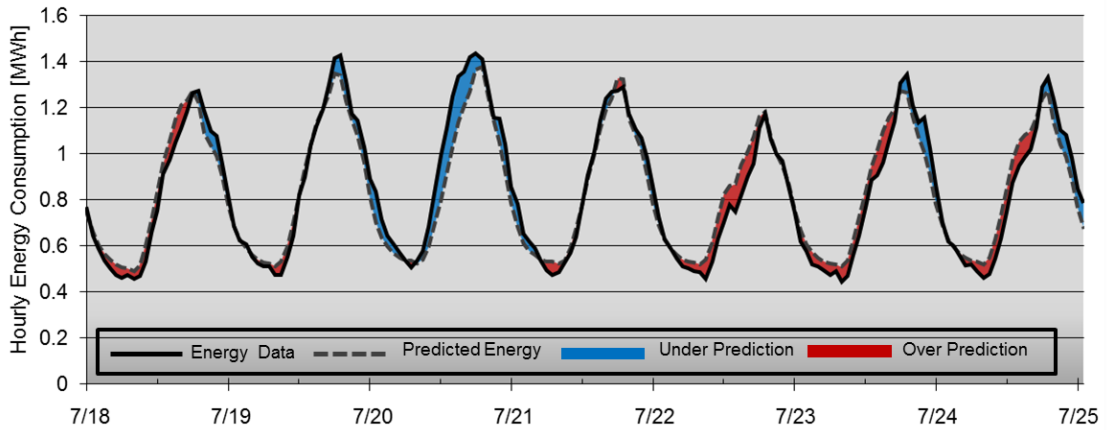
Figure 3.35: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using actual weather data.
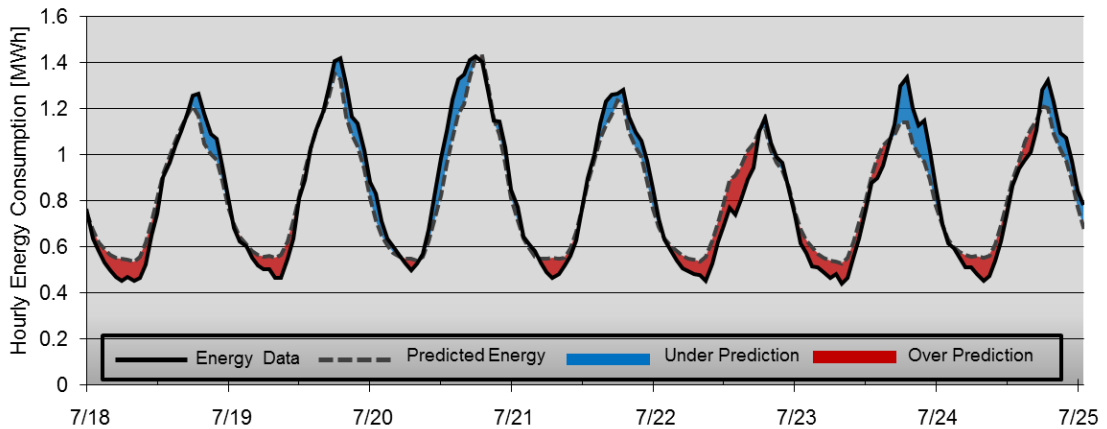
Figure 3.36: Predicted and actual aggregate hourly energy consumption for 250 houses in Houston, TX from July 18 to July 25, 2012 calculated using weather forecast data.

The sensitivity of the smart meter only model to changes in the weather conditions is explored to determine how important an accurate weather forecast is. Figure 3.37 displays the energy consumption predictions using the actual weather data, a constant 1 °C higher outdoor air temperature, a 5 kph higher wind speed, and a 100 W higher solar load. The plot shows both daily and hourly data and the respective sensitivity simulations are plotted from top to bottom.
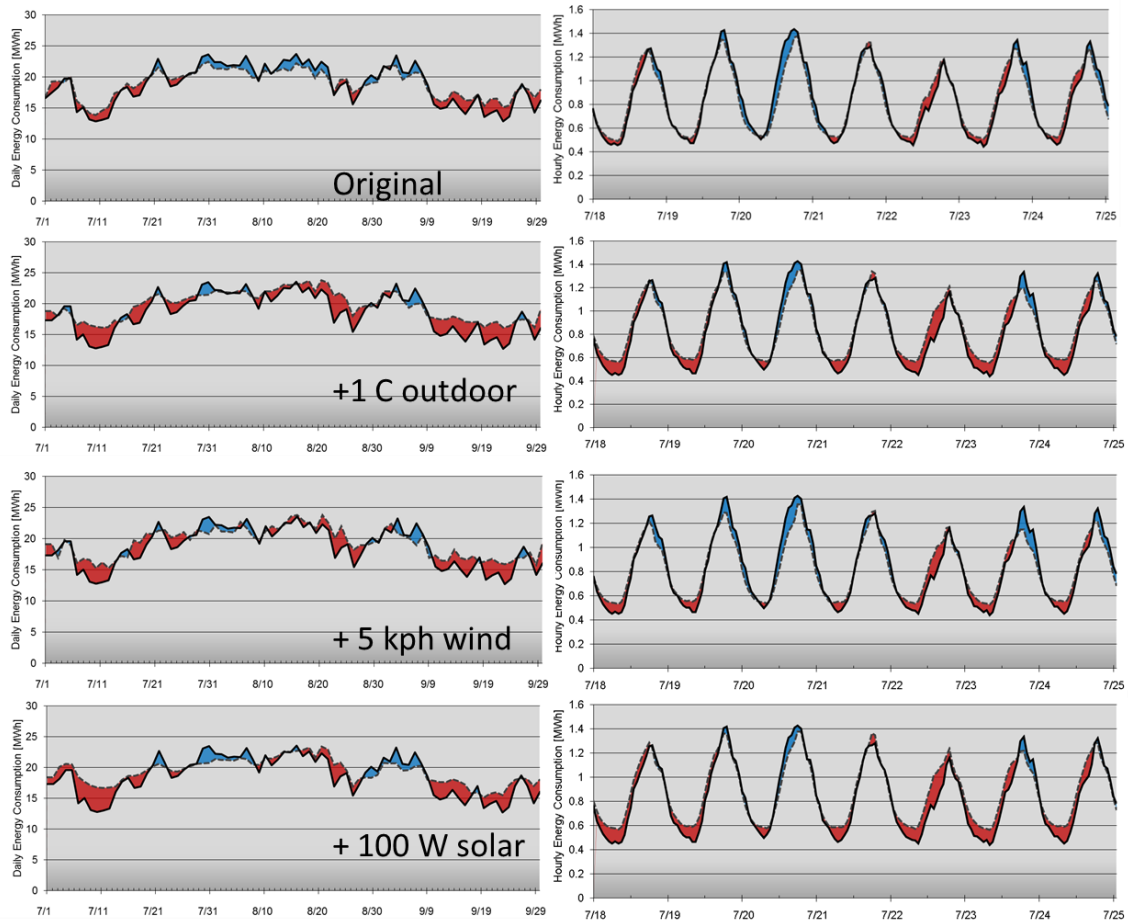
Figure 3.37: Sensitivity of the smart meter only building energy model to weather condition errors. The left plots are the daily energy consumption from July 1 to September 30, 2012 and the plots on the left are hourly consumption from July 18 to July 25, 2012.

84

The RMS error in the predicted daily energy consumption when using the actual weather conditions was 83 MWh and the hourly was 137 MWh. The 1 °C higher outdoor air temperature had RMS errors of 120 and 212 MWh, the 5 kph higher wind speed had RMS errors of 103 and 187 MWh, and the 100 W higher solar load had RMS errors of 115 and 224 MWh.

3.4 Energy Prediction Summary

The performance of the thermostat and smart meter only building energy model are compared to several traditional techniques: cooling degree day, average load, Middle-Eight-of-Ten-Like-Days, and outdoor temperature correlation methods for the 250 homes in Houston, TX from the summer of 2012. Table 3.12 lists the performance metrics of the models using actual weather data. The models are evaluated using the actual weather data to remove any errors associated with the weather forecasts. Table 3.13 lists the metrics when weather forecasts were used as an approximation of how actual predictions would perform. The correlation coefficient is the Pearson product-moment correlation coefficient (R) between the actual data and prediction (Walpole 2007).

Table 3.12: Energy prediction model performance using actual weather data

	Total Error [MWh]	Daily Intervals RMS Error [MWh]	Maximum Daily Error [MWh]	Daily Correlation Coefficient (R)	Hourly Intervals RMS Error [MWh]	Maximum Hourly Error [MWh]	Hourly Correlation Coefficient (R)
Cooling Degree Days	20.3	108	4.19	0.885	457	0.67	0.351
Average Load	0 =NA	220	6.08	0.144	312	0.70	0.676
Middle-8-of-10-Like-Days	10.1	159	6.09	0.594	202	0.52	0.872
Outdoor Temp Correlation	32.4	89	3.39	0.911	227	0.75	0.889
Thermostat Model	8.8	86	3.29	0.942	163	0.66	0.931
Smart Meter Only Model	24.1	83	2.37	0.966	137	0.20	0.952

Table 3.13 Energy prediction model performance using weather forecasts

	Total Error [MWh]	Daily Intervals RMS Error [MWh]	Maximum Daily Error [MWh]	Daily Correlation Coefficient (R)	Hourly Intervals RMS Error [MWh]	Maximum Hourly Error [MWh]	Hourly Correlation Coefficient (R)
Cooling Degree Days	-35.1	135	4.80	0.780	457	0.72	0.332
Average Load	0 =NA	220	6.08	0.144	312	0.70	0.676
Middle-8-of-10-Like-Days	10.1	159	6.09	0.594	202	0.52	0.872
Outdoor Temp Correlation	29.6	102	4.08	0.837	247	0.52	0.854
Thermostat Model	10.2	89	3.72	0.939	191	0.75	0.913
Smart Meter Only Model	24.8	103	3.73	0.889	185	0.51	0.927

When using actual weather data the outdoor temperature correlation method was the best performing traditional method at predicting the daily energy consumption and very near the best at hourly predictions. The thermostat building energy model provided a 3.4 % improvement over the temperature method on the daily energy prediction with an RMS error of 86 MWh from July 1 to September 30, 2012. The correlation coefficient in matching the daily predicted energy to the actual data was 0.942. There was a 28.2% improvement in the hourly prediction RMS error when using weather data at 163 MWh, and the correlation coefficient was 0.931. The thermostat improved the daily prediction by 12.7% with an RMS error of 89 when weather forecasts were used, but it only improved the hourly prediction over the middle 8 of 10 method by 5.4% at 191 MWh RMS error. The correlation coefficient for the hourly prediction using weather data was 0.931 and using weather forecasts was 0.913, where the best traditional methods were only 0.889 and 0.872. The thermostat building energy model predictions saw similar maximum errors to the best traditional technique.

The smart meter only model saw even larger improvements over the temperature method in daily forecasting at 6.7% and hourly at 39.6% with RMS errors of 83 and 137 MWh. The smart meter only model failed to improve the daily prediction when forecast data was used but hourly predictions were improved by 8.4% with an RMS error of 185 MWh. This model did see significant improvement in the maximum error over the best traditional technique. The smart meter only model may outperform the thermostat model predicting general energy consumption, but even with indoor air temperature data it could not predict removable load by controlling thermostats.

Chapter 4: Model Applications

4.1 Introduction

Being able to accurately model buildings and forecast loads with only internet-connected thermostat and/or smart meter data opens the door for many different applications. Future power loads can be estimated given accurate forecasted weather, and if aggregated over the entire utility service territory, they can be used to identify periods of peak capacity. Many utilities are interested in programs that can identify and then alleviate the problems with these peak loads. In the demand response trial and forecasting section of this chapter the results of a DR trial in Houston, TX are presented. This section also investigates the performance of the thermostat building model in predicting the removed load and strategies to improve the capacity with this simulation/forecasting tool. Simply reporting the desegregated HVAC energy consumption data can inspire homeowners to take steps to reduce their consumption. The home efficiency scorecard section of this chapter highlights using the building models to perform a remote energy audit. The thermostat models uses thermodynamic equations to solve for transient indoor air temperatures, thus any indoor temperature control algorithm and weather scenario can be tested. Daily thermostat setpoint schedules are manipulated to reduced HVAC energy usage or operational cost, given variable rate schedules, in the energy efficiency schedules section of this chapter. The computationally light design of the model can perform these tests almost instantaneously, and applications can hold the short attention span of customers. This design also lends itself to large scale simulations for the players in the energy market

aside from DR. Overall, this section presents applications that were developed and tested using the building energy models.

4.2 Demand Response Trial and Forecasting

The building energy model was initially designed exclusively for forecasting and executing direct thermostat control demand response events. Assuming a quality weather forecast and building model, a control simulation can be run prior to the desired event to see how much capacity is available to be removed. Alternatively, given a particular power requirement, simulation forecasts can be performed to determine what control strategy is needed to meet the target. In the summer, load can be removed from a peak period by raising the thermostat setpoint temperature. This causes air conditioners to drift up to a higher operating temperature, reducing the load. Once the prediction method is verified, simulations can be run to test the range of the system under any conditions. In this section thermostat setpoints were changed to remove power during DR events and the building energy model prediction capabilities were tested.

4.2.1 CenterPoint Energy Trial

A test trial spanning from June 15 to September 30, 2012 was performed in the service area of CenterPoint Energy around the Houston, TX area. Customers were acquired almost entirely by digital (email) marketing, social networks, and a limited amount of press/media attention. This was executed through emails to City of Houston employees, CenterPoint employees, contractors who perform energy efficiency installations under CenterPoint programs, and those customers that had an

eligible thermostat already installed. Customers voluntarily signed up for the pilot in exchange for energy efficiency tools (see sections 4.3 and 4.4) and a $25 Home Depot gift card. In total, 160 houses (unique smart meters) with 206 thermostats participated in the trial, which according to Adib (Pioneer Energy Consulting 2012) would be adequate to achieve a reliable figure representing the average deemed load reduction savings. Adib stated "a random sample of 100 to 170 pilot participants would result in a probability of 95% that the estimated average deemed savings per thermostat for the e5 Program Pilot in CenterPoint Energy service area will fall within 3% of the true but unknown average for the relevant population (potential adopters of this new technology)."

Funding for the trial came from leveraging the amount of removable load during an actual grid emergency event in CenterPoint's 1-hour alert Emergency Response Program and their R&D budget. No actual emergency events were called over the summer, but CenterPoint did initiate one test event. Small subsets of the thermostat population were tested several times (June 28 and August 2, 2012), but only the major events (August 15, 27, September 6, 7, and 13, 2012) are documented in this dissertation.

The demand response events tested had three main phases: alert/precooling, period of removed load, and recovery. An alert would trigger the system to generate new thermostat setpoint schedules for all the participating thermostats. Because the energy model was built for computational speed, producing, sending, and receiving the schedules usually took less than 5 minutes for the entire set of targeted thermostats. A small subset of the population required 5 to 15 minutes or did not

receive the schedule at all due to internet connection issues. The rest of this time preceding the event was used for lowering the indoor temperature as low as possible, necessary, or allowed by the customer (2 °C below their lowest daily setpoint). This precooling period is designed to extend the length of time in the actual event when the indoor air temperature can drift up to the event setpoint. A 1 °C setpoint increase above the user's thermostat setpoint schedule was tested during the period of shifted load. After the event concluded the system would enter the recovery phase where the thermostat setpoint was returned to its original scheduled value. This recovery phases resulted in the majority of air conditioners needing to turn ON and added a new spike in the load.
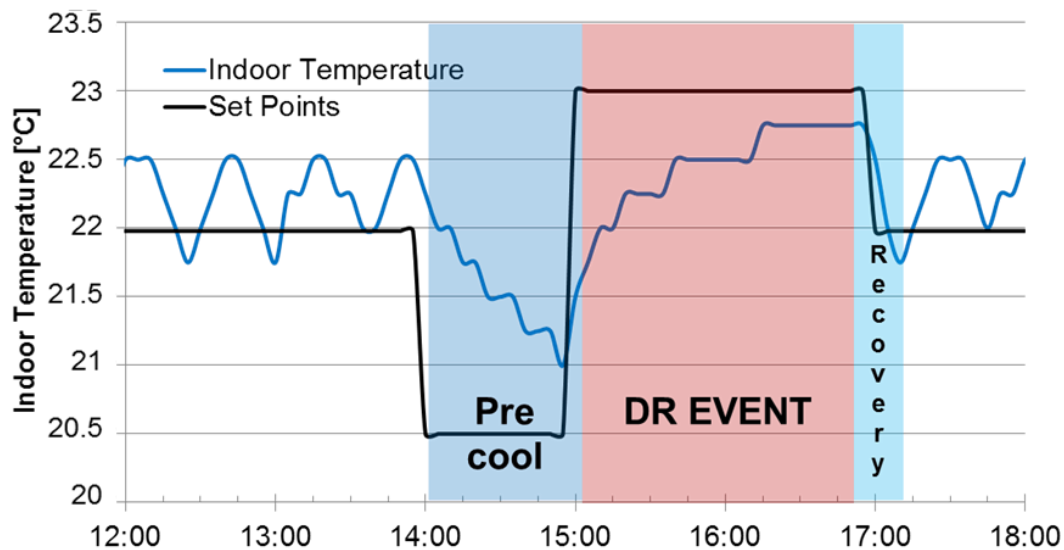
Figure 4.1: Indoor temperature and thermostat setpoint profile of a house in Houston, TX during a demand response event on August 15, 2012. This plot illustrates the various phases of a thermostat DR event.

During the event customers had the ability to "opt out" and remove themselves

from participation. They were only able to do this through the web portal on the first event, but were able to do it at the thermostat or mobile app for the subsequent events. If the proper opt out procedure was not performed adjusting the thermostat setpoint would only be applied temporarily until the server would reinstitute the demand response schedule (roughly 5-15 minutes).

The Middle 8-of-10 Preceding Like Days Model was employed from ERCOT's Emergency Interruptible Load Service Default Baseline Methodologies Manual to baseline the load. This method assumes that the energy consumption of recent days would most likely reflect what would occur on the day of an event. Ten days are selected preceding the event that fall into the same category. All the tests were performed on weekdays so any weekend, holiday, or past demand response event days were excluded and an earlier weekday was selected until there were ten. The lowest and highest daily energy consumption sum was removed to produce the middle eight days, and a time average of this became the baseline. This analysis is performed for each individual house, and aggregated for the entire participating set.

The first full test event occurred on August 15, 2012 from 15:00 to 17:00 CDT. An alert was simulated at 14:00 by sending demand response schedules to the thermostats. The test included 140 homes with 187 thermostats, 10 failed due to internet connection issues. Three users managed to opt out of the event and 34 attempted to by adjusting their thermostat. Figure 4.2 plots the baseline power (black) for the 140 smart meters involved in the event and the actual power consumed on the demand response day (orange). The difference between the two power signals is given in green. The targeted stations that did not receive a demand response schedule

92

and users who opted out or adjusted the thermostat were not excluded from this aggregation because these cases are inherent to the way a real event would occur, and need to be included. Note the power spike between 14:00 and 15:00 for the precool, power drop right at 15:00 and slow increase to 17:00, then the later spike after the event ends at 17:00.
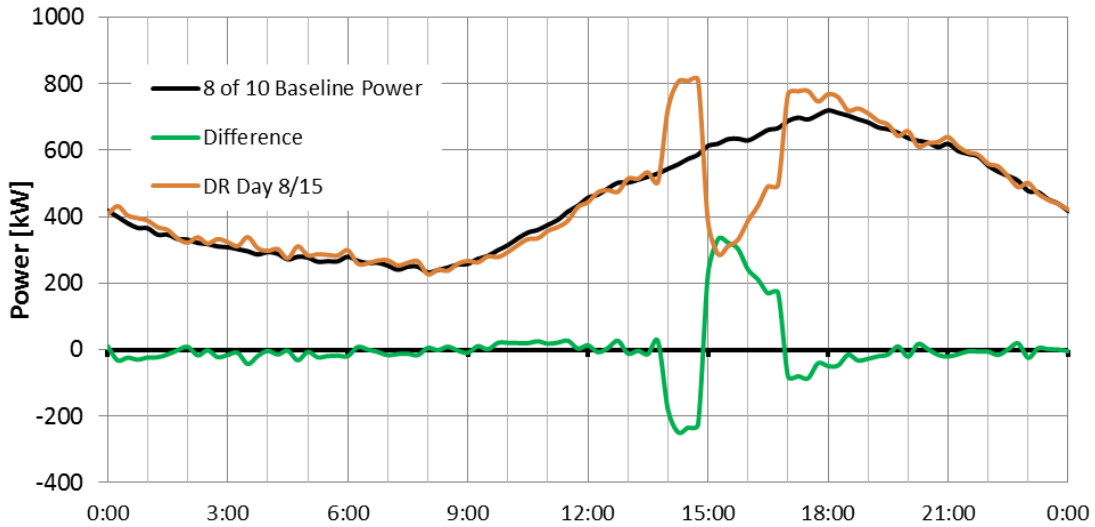
Figure 4.2: August 15, 2012 demand response event 15:00-17:00 using 187 thermostats.

Figure 4.3 plots the HVAC power data and indoor temperature averaged over the entire set from the event, as well as the average of all the building model predictions for that set. Thermostats were removed from this analysis if they did not behave in a predictable manner (opted out or adjusted the thermostat more than once). The temperature model did not respond as abruptly as the actual data, and it over-predicted the average temperature for the entire duration following the event. The model also under-predicted the power consumption during the event.

Figure 4.3: August 15, 2012 demand response event 15:00-17:00 average house data and prediction.

The model predicted that the average HVAC power during the 2 hour event and 2 hours both preceding and following the event would be 1.95 kW. The actual computed average power was 2.02 kW which was a 3.6% error. The 5-minute RMS error in the difference between the average predicted and computed energy consumption was 1.1 kWh.

A test event of this nature also occurred on September 6, 2012 but ran from 16:00 to 18:00 CDT. Figure 4.4: September 6, 2012 demand response event from 16:00-18:00 using 200 thermostats plots the baseline and total power throughout the day and Figure 4.5: September 6, 2012 demand response event from 16:00-18:00 predictions and data the average predicted power and indoor air temperature compared to what was computed and measured.
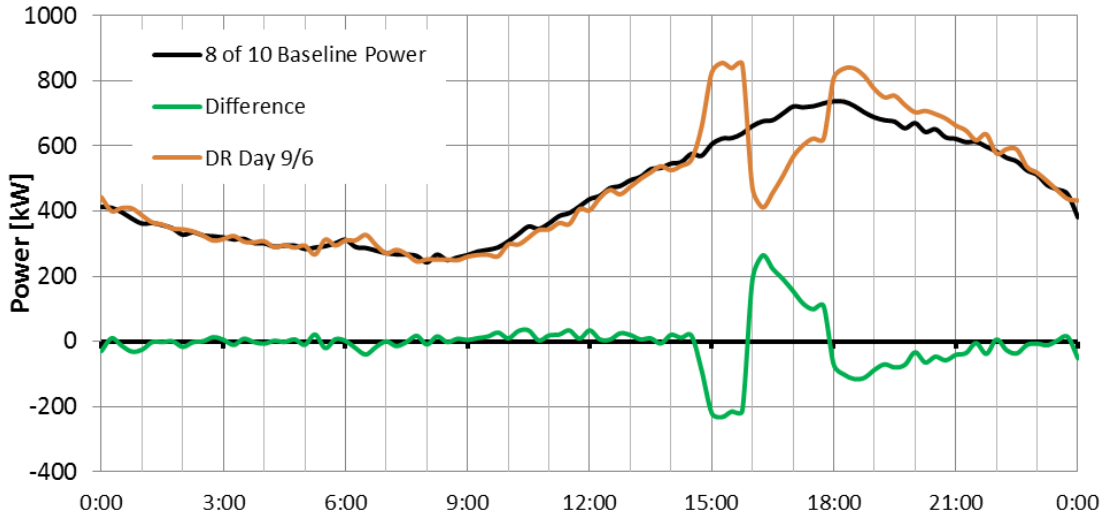
Figure 4.4: September 6, 2012 demand response event from 16:00-18:00 using 200 thermostats.

Figure 4.5: September 6, 2012 demand response event from 16:00-18:00 predictions and data.

The model predicted that the average HVAC power during the 2 hour event and

2 hours both preceding and following the event would be 1.71 kW. The actual computed average power was 1.79 kW which was a 5.3% error. The RMS error in the difference between the average predicted and computed energy consumption was 1.6 kWh.

A DR test was initiated by CenterPoint on August 27, 2012 starting at 15:30 with only 50 minutes of alert time and 2.5 hours of shifted power. Figure 4.6: plots the baseline and total power throughout the day and Figure 4.7 plots the average predicted power and indoor air temperature compared to what was computed and measured.
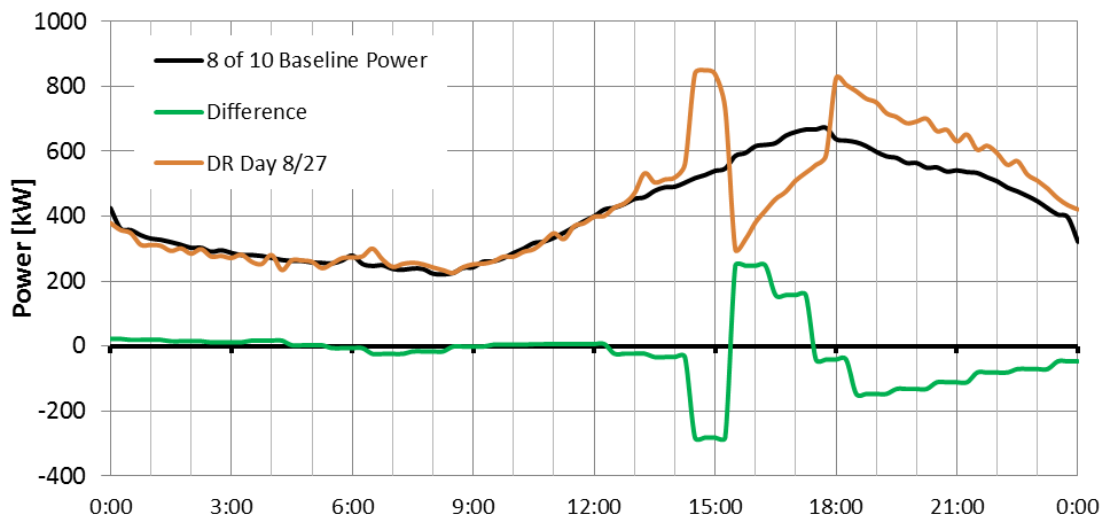
Figure 4.6: August 27, 2012 Demand Response Event 15:30-18:00 using 186 thermostats.
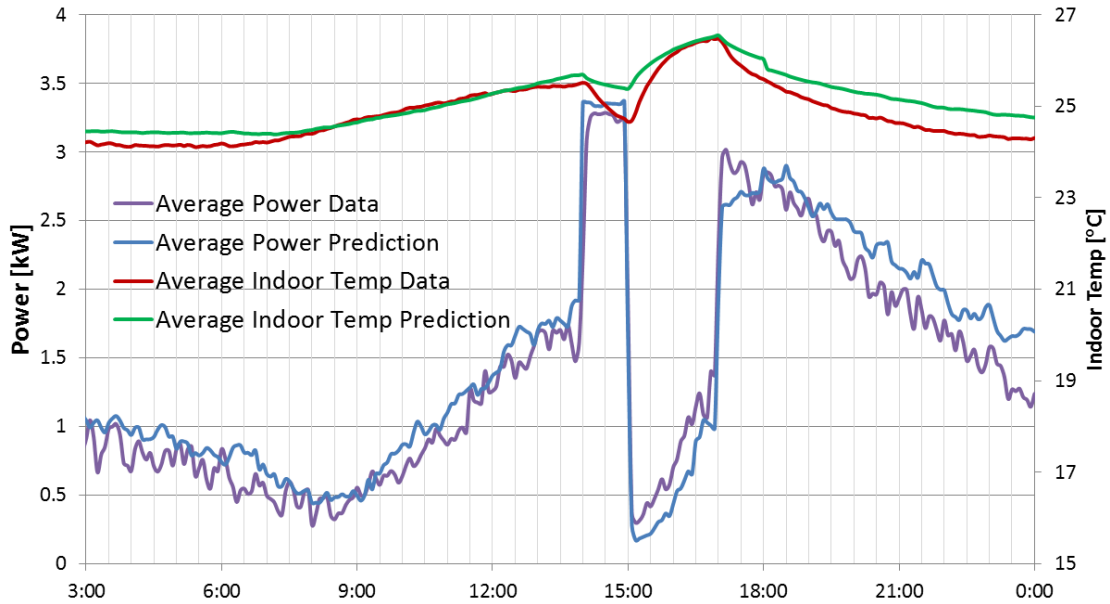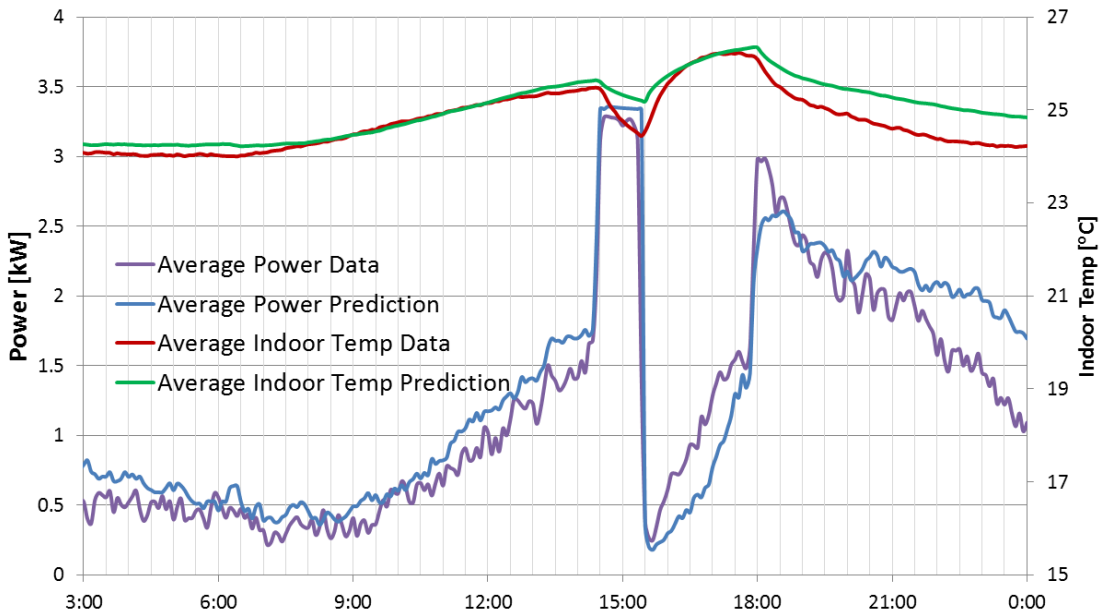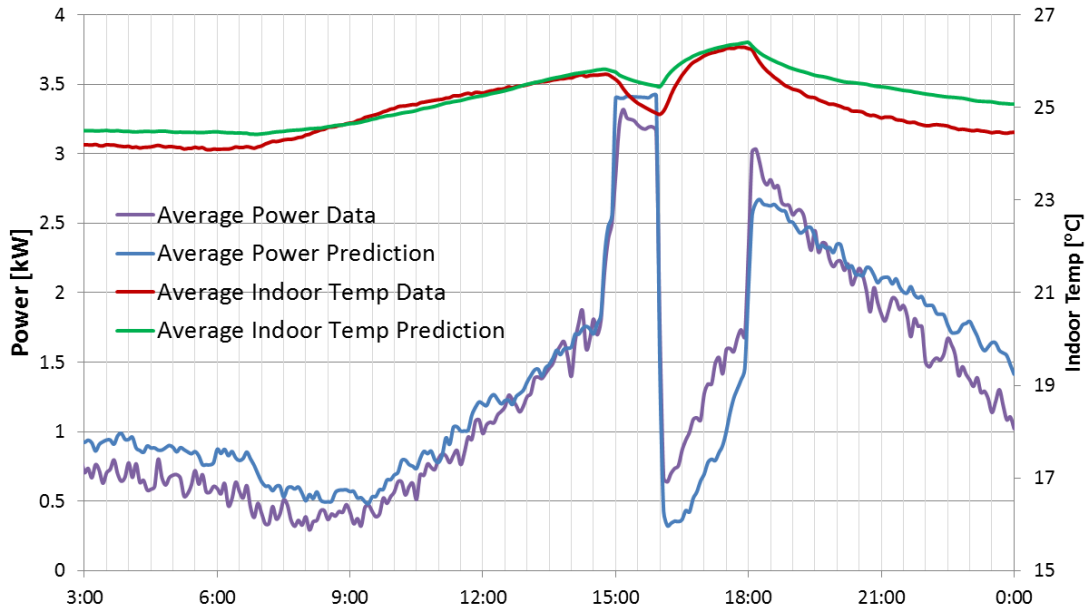
Figure 4.7: August 15, 2012 Demand Response Event 15:30-18:00 average house data and prediction.

The model predicted that the average HVAC power during the 2 hour event and 2 hours both preceding and following the event would be 1.83 kW. The actual computed average power was 1.99k W which was a 8.9% error. The RMS error in the difference between the average predicted and computed energy consumption was 1.5 kWh.

Table 4.1 summarizes the results of these three DR tests events. The total per thermostat average for the three tests over the entire event duration was 1.13 kW. This is about one third the power of the average modeled HVAC unit. This number is expected because roughly half the thermostats did not make it through the +2 hour events without needing to turn on. About half of those did not make it because the users schedule was set to change to a setpoint low enough to be effected. The average RMS between the indoor air temperature and predicted model air

Table 4.1: Summary of the 3 1-hour alert demand response tests

DR Test	15-Aug	27-Aug	6-Sep
Total Shifted Power [kW]	247.35	229.19	171.58
Number of Thermostats	187	186	200
Normalized Shifted Power [kW]	1.32	1.23	0.86
Opt Outs (of event) [%]	1.60	13.98	14.50
Adjusted Thermostat [%]	18.18	19.35	19.00
Error in Load Prediction [%]	3.58	5.25	8.89
RMS Error in Energy consumption [kWh]	1.11	1.63	1.49

temperature was 0.33 °C which is better than the averages over the entire day. The average power predicted by the model during the demand response event was only 0.34 kW per thermostat. The difference in these numbers is due to the forecast model not knowing if a station will opt out and apply a different schedule. This energy error was consistent in every test and an opt out waiting factor can be applied to reach a stronger agreement. Simulations were also run on four other days in the summer (July 7 and 30, August 16 and 28, 2012) to determine the removable load in this late afternoon timeframe. The simulations started on a measured average outdoor air temperature (around 28.6-30.1 °C), but were then rerun given outdoor air temperatures of lower and higher averages (18.5-40 °C). These were used to determine a load curve for the average summer day. Given a daily average air temperature (which could be equivalently expressed in cooling degree days) the curve can be used to determine the approximate energy available in the late afternoon. The three main events and two smaller test events were plotted to test if the curve was accurate. This curve is plotted in Figure 4.8.
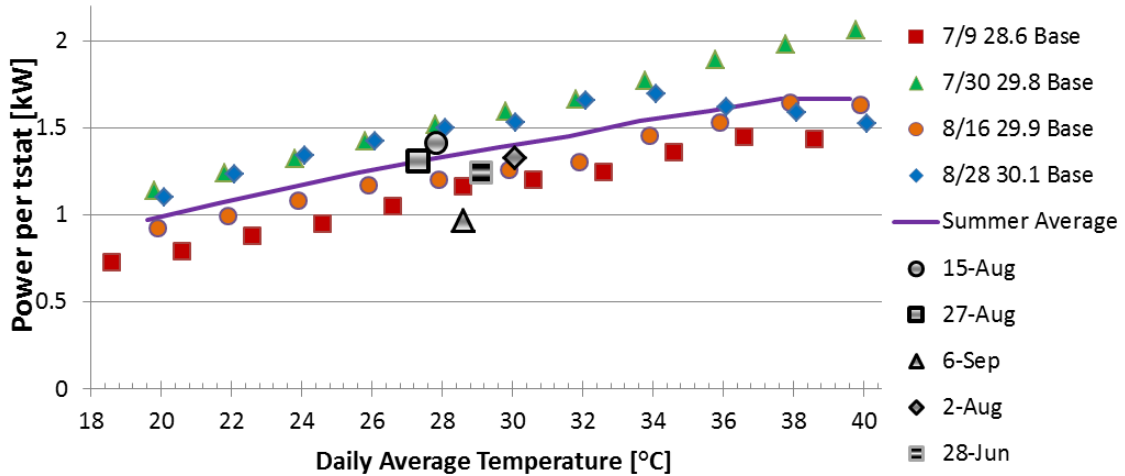
Figure 4.8: Removable power load curve per thermostat verses daily average temperature given 1 hour alert.

A test demand response event was performed on September 7, 2012 so see if the process could be improved. Three treatment groups of 60, 59, and 59 thermostats were created. Each group had roughly the same number of houses with multiple thermostats and ones that had experienced connection issues in previous events. The first treatment group, labeled adaptive precooling, was given three hours to prepare for the upcoming event. An algorithm was developed using the thermostat building energy model to determine the amount of precooling needed to have the minimum amount of load in the DR event. This algorithm iteratively increased the duration of the precooling setpoint preceding the start of the event until it predicted the AC would not turn ON during the event or further precooling would be uncomfortable. This algorithm was used to determine the precooling schedules for the adaptive precooling treatment group.

The second treatment group, labeled standard precooling, operated like the events test prior with only 1 hour of alert / precooling. The third group, labeled no

99

precooling, was simulated to receive a demand response alert right at the start of the event. These schedules were immediately set to the higher demand response thermostat setpoint temperature. Figures of the power and temperature can be viewed for each test event in the appendix. Figure 4.9 summarizes the phases of the three test demand response strategies.
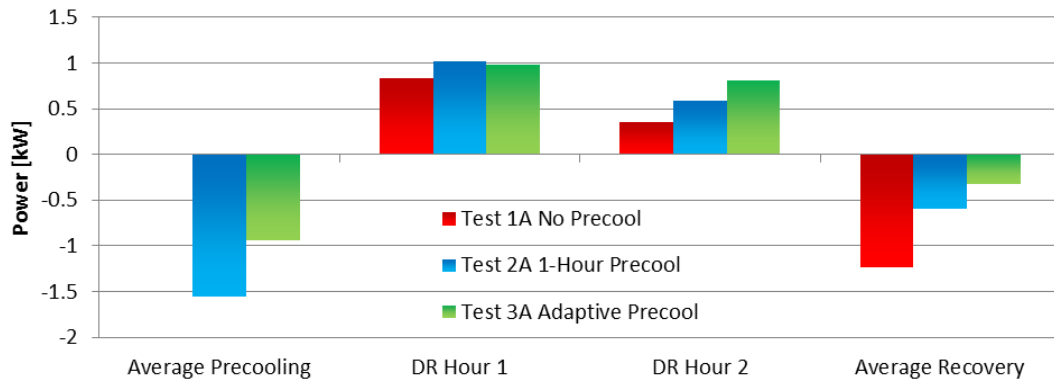
Figure 4.9: Power summary of the three treatment groups on September 7, 2012 demand response test event. Values are averages per phase of the event per thermostat.

The adaptive precooling averaged lower precooling power than the 1-hour precooling because it was implemented over 3 hours and not just 1 but required more overall energy. Both precooling strategies outperformed the no precooling strategy during the event by shifting more power. This was expected because the no precooling strategy would spend less time drifting up in temperature and more time cycling to maintain it than the precooling strategies. The no precooling was significantly worse in the recovery phase as well because more homes were at the higher demand response temperature exiting the event. Figure 4.10 plots the average indoor air temperature of each of the three strategies. The adaptive precooling

maintained the lowest average indoor air temperature during the event and recovery while the no precooling strategy was significantly higher in temperature than the other strategies. This was expected because the precooling is performed to minimize this discomfort while still yielding a significant amount of shifted load.
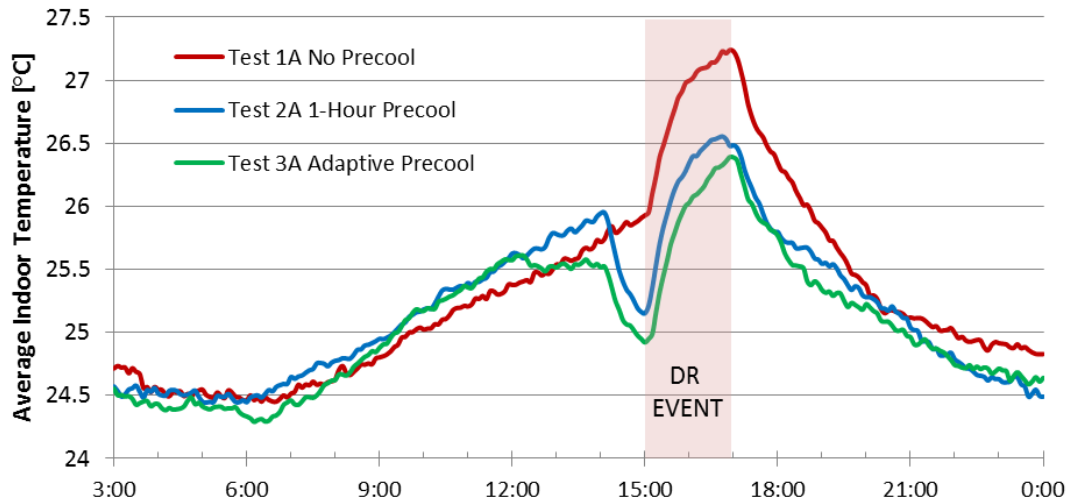
Figure 4.10: Average DR treatment group indoor temperature on September 7, 2012 demand response test per thermostat.

4.2.2 ERCOT Trial

Another set of 324 thermostats was included in the ERCOT Emergency Interruptible Load Service program spanning all of Texas. This demand response program only had 30 minute alerts and had different day periods to call an event (9:00-14:00, 14:00-17:00, 17:00-21:00, and all times else). Only one test event was called by ERCOT on September 13, 2012 from 10:30 to 11:30 CDT. The system only had access to the 101 CenterPoint service area thermostats so the analysis could only be performed with those. The results of this test are plotted in Figure 4.11 along with load curves generated from simulations of test events beginning in each of the time

intervals.
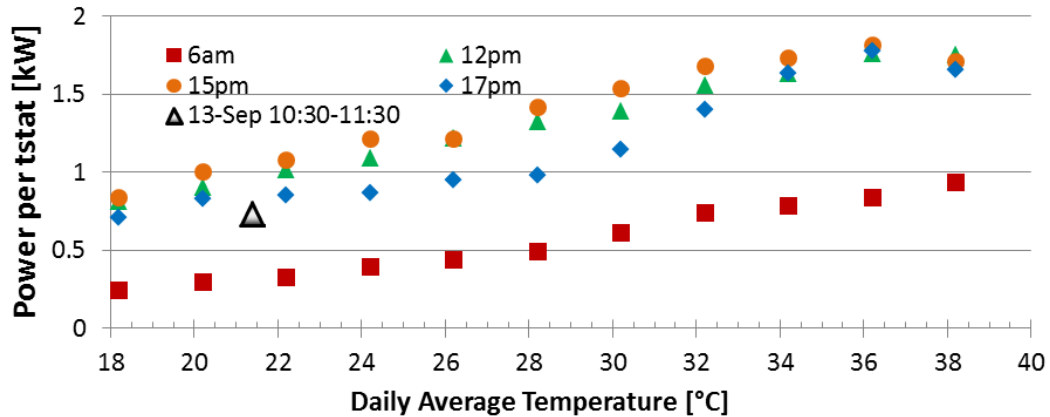
Figure 4.11: Removable power per thermostat verses the daily average temperature for each of the time intervals. The ERCOT initiated test on September 13, 2012 is also plotted using a subset of the population.

The test events in the two pilot programs validate the demand response to a certain extent. One weakness was that all the events in the CenterPoint trial occurred on similar weather days. This happened because each major test was trying to determine the maximum total load that can be shifted using this method and warmer-than-average days were always chosen to test that. But note that these were not significantly different. There was not much variation in the daily average outdoor temperature over the entire trial.

4.3 Energy Efficiency Schedules

A service was developed to improve the efficiency of thermostat setpoint schedules for customers who participated in the CenterPoint DR trial and the Earth Networks e5 program. The service used the building energy models and weather forecasts to test techniques that could reduce the energy consumption of operating the

AC. Customer thermostat setpoint schedules are adjusted using three mechanisms: smart setback, setpoint smoothing, and precooling (or preheating). Smart setback gives customers the ability to set what temperature they want it to be at a particular time, as opposed to just the setpoint. Setpoint smoothing shaves off HVAC cycling before a setback period, and precooling takes advantage of higher efficiency operating conditions. Customers were able to start utilizing this scheduling service in September 2012, but the actual energy saving they produced has yet to be determined.

4.3.1 Smart Setback

Raising the thermostat setpoint temperature during the summer or reducing them in the winter is known to reduce energy consumption. Maintaining these elevated and reduced setpoints requires less energy because the indoor conditions are closer to a thermodynamic equilibrium with the outdoor environment. Figure 4.12 simulates controlling a house at 22 °C and 26 °C where the higher setpoint consumes 29% less energy.
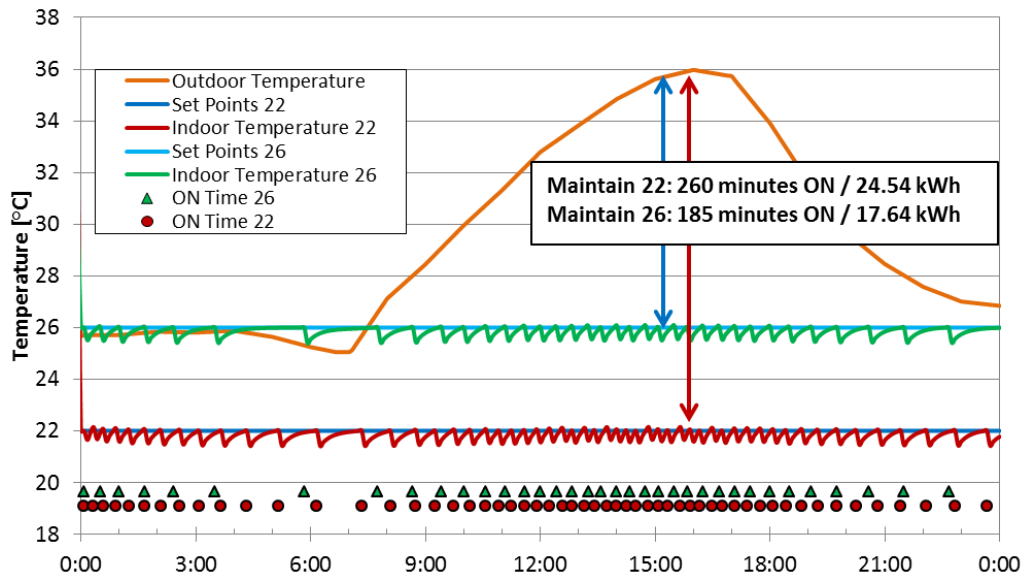
Figure 4.12: Setpoint simulation.

The problem with maintaining these setpoints is that they are uncomfortable. Setback is a technique homeowners have developed to acquire some of the benefit of uncomfortable setpoints while the home is unoccupied. Figure 4.13(a) simulates a house with constant 22 °C setpoints and Figure 4.13(b) with a 3.5 °C setback from 9:00 to 17:00. The energy consumption is reduced 11%.
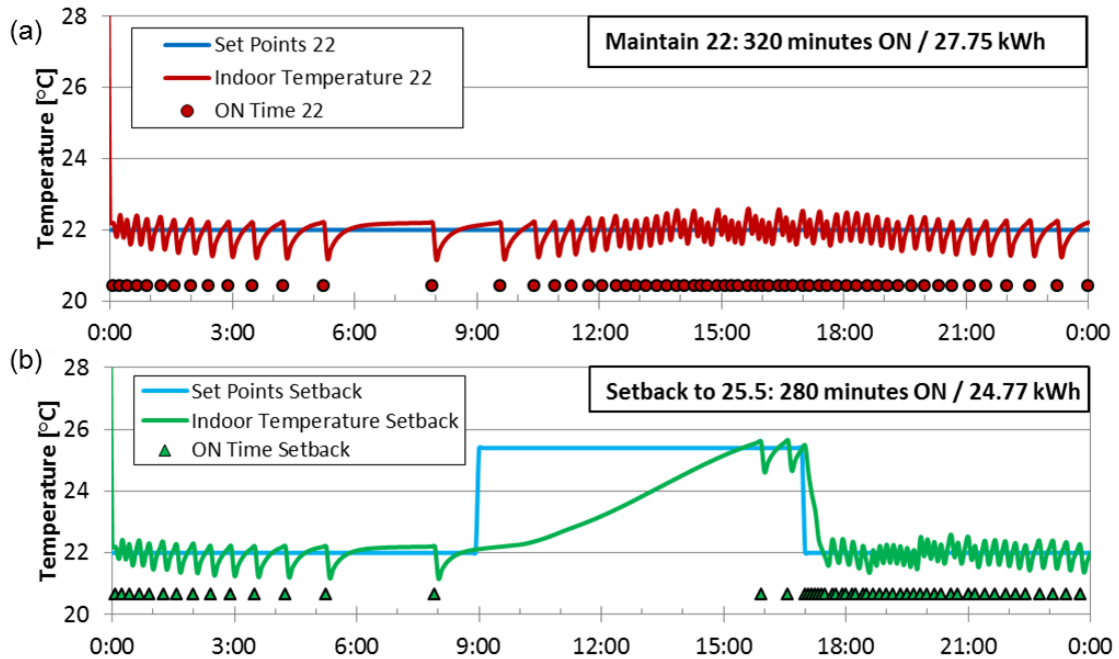
Figure 4.13: Setback simulation. (a) Maintaining a setpoint of 22°C. (b) Setback of 3.5 °C from 9:00 to 17:00.

Smart setback works by calculating the time required to shorten a setback setpoint in order to bring the actual indoor air temperature to the desired temperature. The algorithm simulates the indoor air temperature with the building model and iteratively shortens the setback period until the desired temperature is forecasted to be met. In the following example (Figure 4.14) a user sets back to 27 °C while away at work then has it programmed to change to 23 °C when they return at 18:00. A traditional thermostat will wait until that programed time (18:00) to actually turn ON,

and the user may come home to an uncomfortable house.
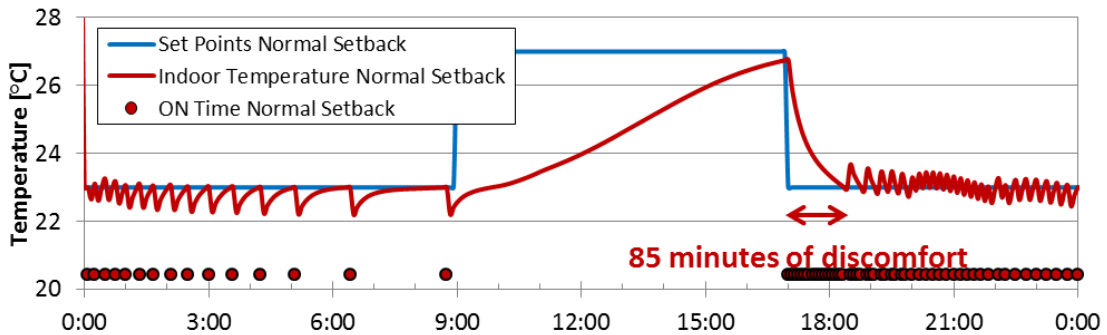
Figure 4.14: Simulation of traditional setback resulting in 85 minutes of discomfort.

If 18:00 is programmed with smart setback the algorithm will determine how far in advance that 23 °C setpoint needs to be implemented so it will be 23 °C at 18:00 (Figure 4.15).
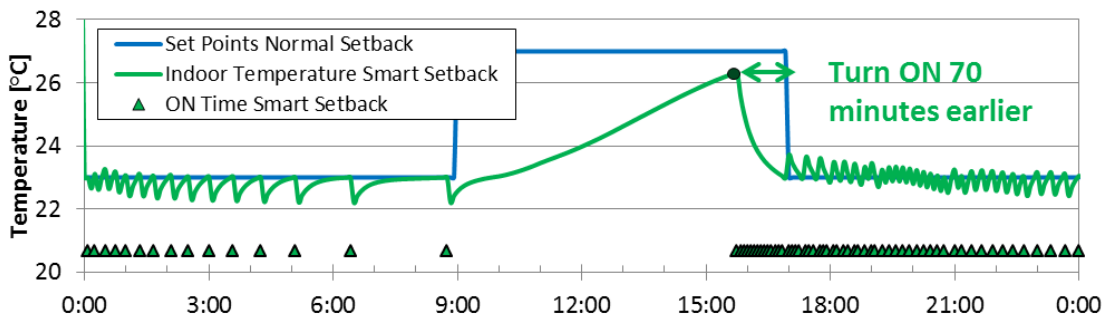
Figure 4.15: Smart setback begins the comfortable setpoint 75 minutes earlier.

In this particular example the smart setback schedule actually would consume more energy than the traditional setback schedule because less time is spent at the higher temperature. Overall, this technique saves energy because it allows the user to be more aggressive with their setbacks. It converts users who originally held constant temperatures because of fear of being uncomfortable into ones wanting to utilize the full energy saving potential of setback. Figure 4.16 shows that the majority of users in Houston, TX set back 2 °C or fewer, so there is significant potential to utilize this

105

technique.
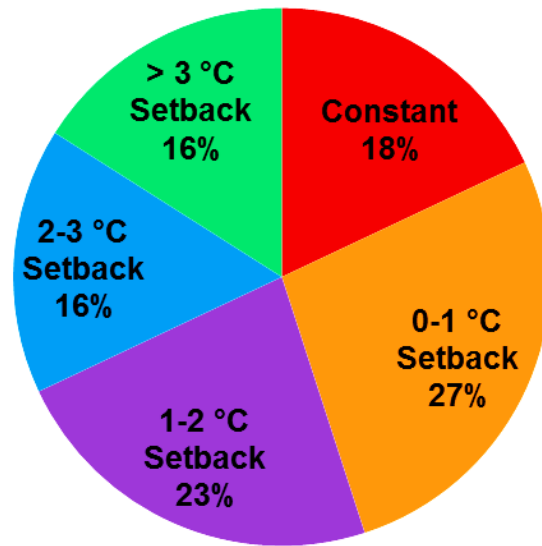
Figure 4.16 The percentage of customer setbacks from 250 houses in Houston, TX from the summer of 2012

It is also shown that improperly guessing this shortening time can result in discomfort or even wasted energy. In the following example it was determined that 70 minutes earlier is enough time to adjust the return setpoint to achieve the desired 23 °C at 18:00 on the average day. Figure 4.17(a) simulates how the original 70 minutes is not enough on hotter days and results in 115 minutes of discomfort. Figure 4.17(b) shows how 156 minutes were needed on this day. Figure 4.18 depicts the opposite case where the 70 minutes was too much on a cool day, and 25 minutes of unnecessary ON time was incurred (3.29 kWh for this house).
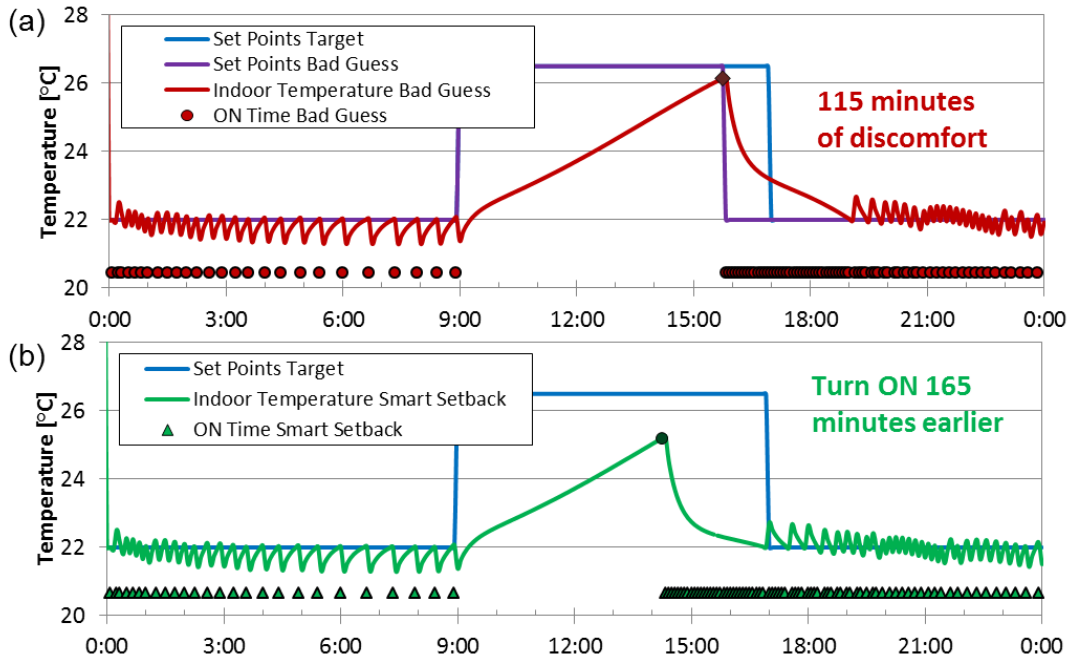
Figure 4.17: Smart setback simulation on a hot day. (a) Bad guess of only 70 minutes to return the setpoint. (b) Smart setback 165 minutes earlier.
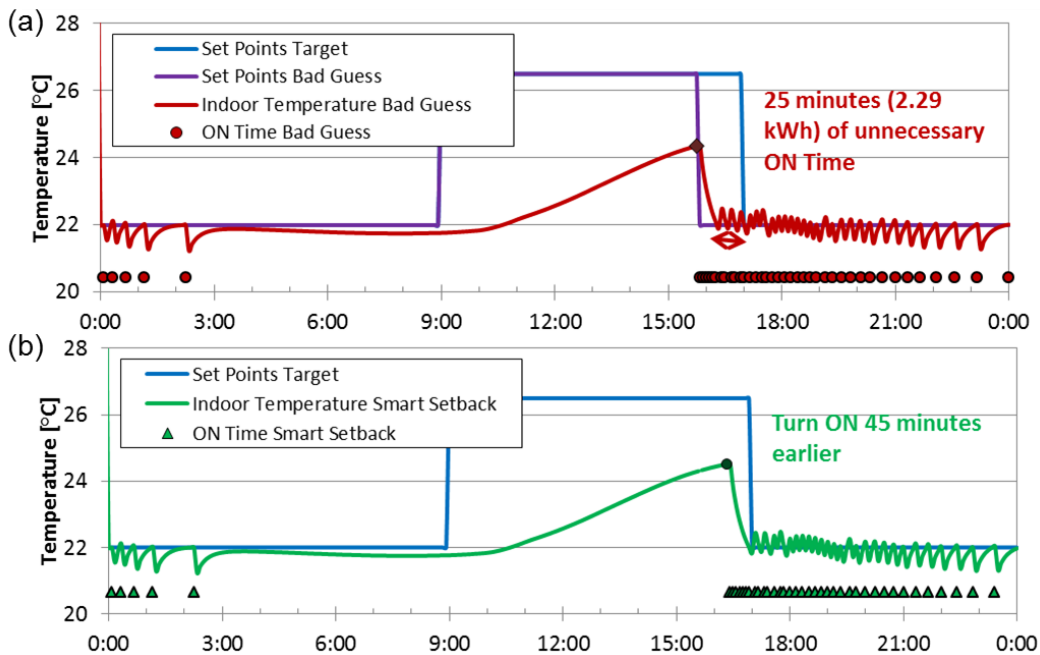
Figure 4.18: Smart setback simulation on a cool day. (a) Bad guess of only 70 minutes to return the setpoint. (b) Smart setback 165 minutes earlier.

Figure 4.19 displays an example of a thermostat schedule generated for a house at midnight where the smart setback predicted the 6:00 setpoint needed to be started at 3:15 and the 18:00 at 16:50. The red indoor temperature high threshold line shows the maximum temperature the customer wanted throughout the day, and the smart setback schedule was produced to achieve that with the minimal amount of energy consumption.
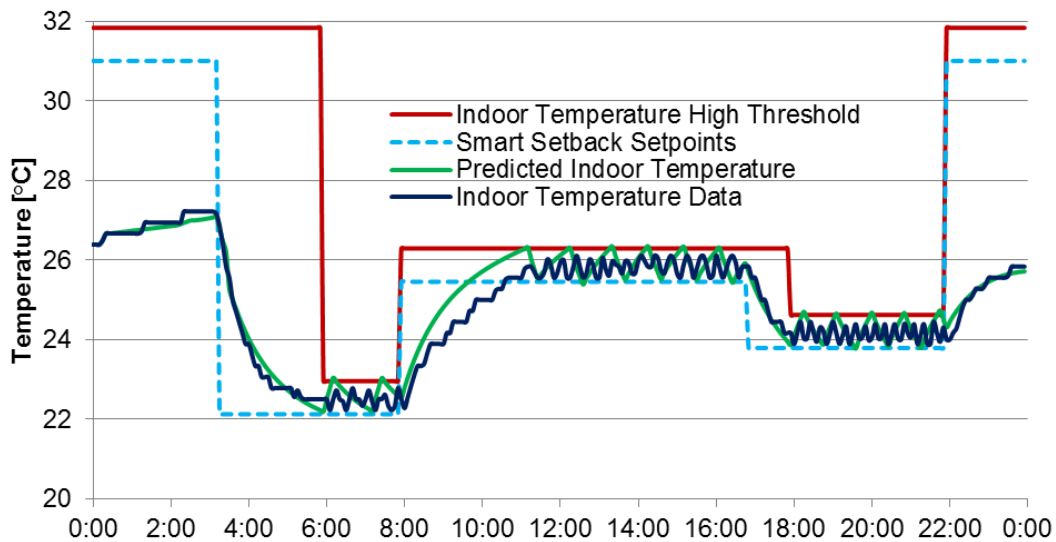
Figure 4.19: Model predicted schedule with smart setback for 6:00 and 18:00 plotted with the actual data recorded by the thermostat for a house in Mt. Airy, MD on August 17, 2012.

4.3.2 Setpoint Smoothing

Another mechanism for energy savings is setpoint smoothing. This occurs in the hour before a schedule is nearing a jump up to a high setback temperature. The setback temperature is started earlier if the temperature will not rise more than 1 °C by the end of the current period. This usually shaves off one or two duty cycles from the lower, less efficient setpoint temperature. Figure 4.20 shows how setpoint

smoothing removes cycles before the 8:00 and 22:00 setbacks.
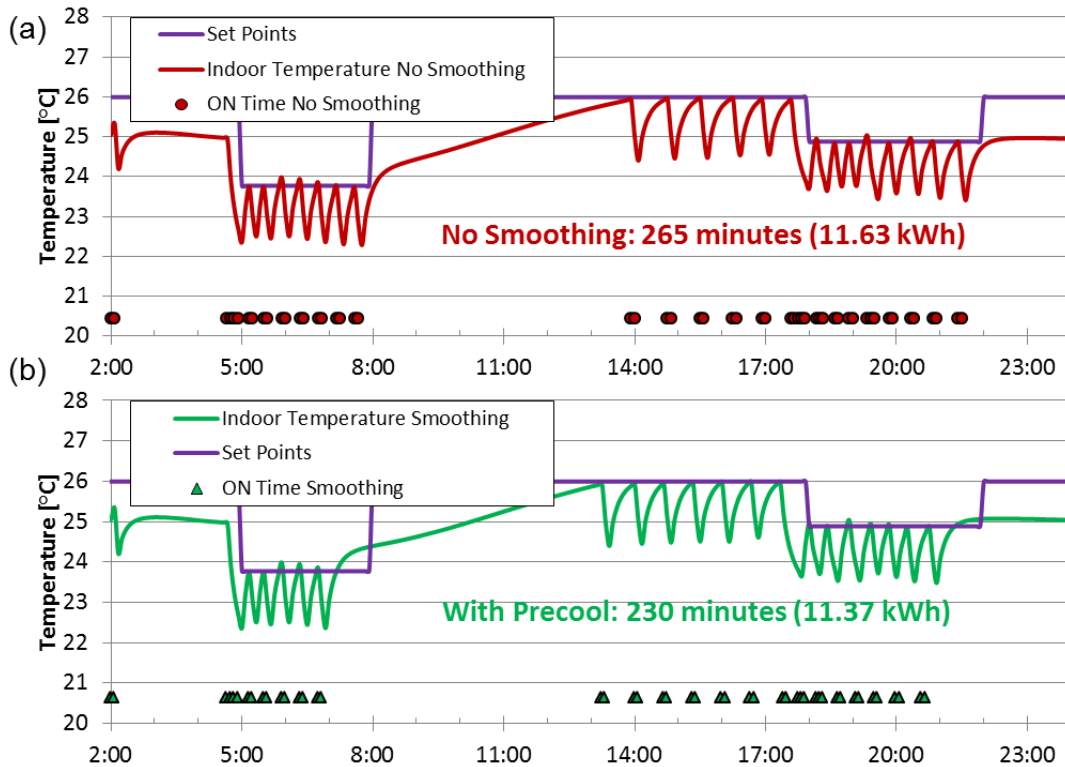
Figure 4.20: Setpoint smoothing simulation. (a) without setpoint smoothing. (b) With setpoint smoothing.

4.3.3 Precooling

The third mechanism involves precooling in the early morning (when an air conditioner gets the most cooling per input kW of power) and/or prior to a price increase. The algorithm iteratively finds the balancing point between utilizing efficiency and precooling too much by abandoning the testing of additional precooling when the overall consumption begins to increase. Figure 4.21 simulates how precooling at the more efficient 2:15 reduces the amount of cooling required at the less efficient 18:00.
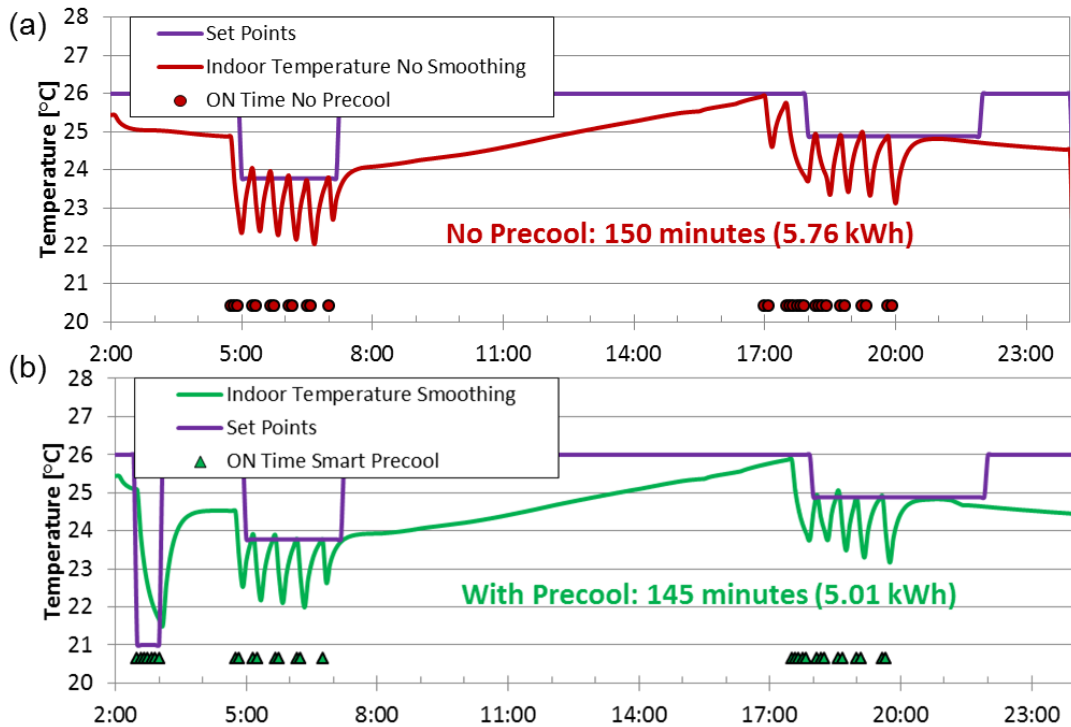
Figure 4.21: Precooling simulation. (a) No precooling. (b) 25 minutes of precooling at 2:15 to reduce the energy consumption in the late afternoon.

The scheduling service simulates schedules created with combinations of the three of these mechanisms and chooses the schedule that is predicted to consume the least amount of energy. The same analysis can be performed if a price schedule is introduced, especially if there is a variable rate. Precooling before a price increase can reduce the cost of operating at that time. These mechanisms were developed to demonstrate how the building energy models are be used to improve the energy efficiency of homes. While a complete study of their effectiveness has not yet been performed, a trial is scheduled to be performed in the summer of 2013. More on this is presented in Chapter 5.

4.4 Home Efficiency Scorecard

Smart power meters give homeowners the opportunity to investigate their energy consumption in more detail than traditional metering. However, most people are not educated on the subject and cannot take advantage of this new technology. This section outlines a home efficiency scorecard that was developed to help educate customers about their energy consumption. The building energy models are used to disaggregate the loads and test the house under various weather conditions to determine performance rankings. These rankings are used to target what are the most effective improvements customers should make.

Companies such as OPower have been stepping in to bridge this education gap by working with many utilities nationwide to provide energy consumption reports (Smart Power 2013). These reports compare the customer's energy consumption to similar sized homes in their neighborhood, as well as offers tips to reduce energy consumption (Home Energy Report 2013).
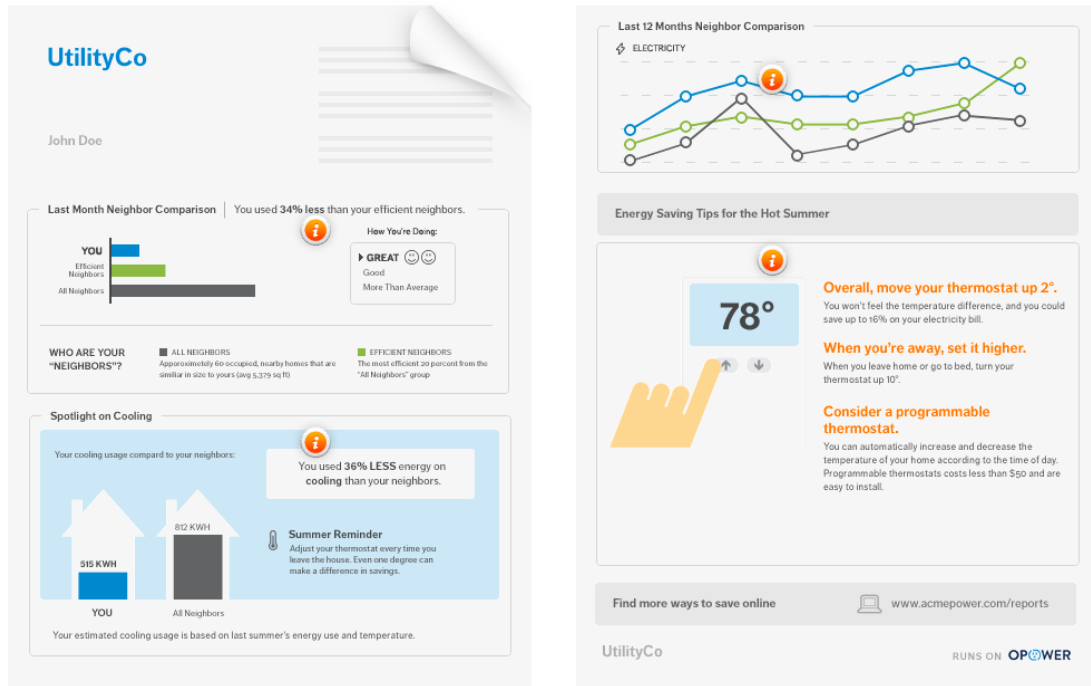
Figure 4.22: OPower energy report (Home Energy Report 2013).

At the time this dissertation was completed, OPower has claimed their reports have driven customers to reduce their energy consumption by over 2 TWh. OPower has achieved this by offering very little customized information to the customer. The building energy models generated from thermostat and smart meter data are tools that can be used to investigate home energy consumption further.

4.4.1 Thermostat Building Energy Model Scorecard

This section highlights a home efficiency scorecard distributed to customers with internet-connected thermostats. Example scorecards are shown in Figure 4.23 and Figure 6.1, Figure 6.2, and Figure 6.3 in the Appendix. Another method is also explained to further disaggregate base and appliance use from total energy consumption using the heat transfer metric from Section 3.3.1.

112

Figure 4.23: Home Efficiency Scorecard for a customer in the CenterPoint Energy utility rating consumption from June to September 2012.

The energy consumed by the HVAC system is disaggregated from the total consumption using the internet-connected thermostat data. The building model is

used to test how much energy would have been consumed if an alternative thermostat strategy was implemented. In these reports the measured cooling energy (calculated using the measured ON time data and power curve) is compared to simulations using the energy efficiency service from section 4.3 and operating with Energy Star setpoints (Thermostat Guidelines 2013). If a customer utilized the energy efficiency service their consumption is compared simulations that did not utilize these setpoints.

The thermodynamic models allow the HVAC consumption to be further divided into the individual sources. Heat is removed or added from the indoor environment by the HVAC system in response to energy added or removed via three major sources: conduction through the envelope, infiltration, and solar loading. Energy in the HVAC system can also be invested in the thermal mass of the internal air or structure as well, but overall the HVAC is designed to combat these main heat sources and sinks. The components of the total energy consumption of each source are estimated by running simulations on past data with key parameters removed. In the case of estimating how much energy is lost due to infiltration a simulation is run where the infiltration coefficient is zero. The control algorithm chooses to turn the HVAC ON and OFF to match the previous setpoints, and if infiltration is significant it will predict a different overall runtime than what actually occurred. This difference is thought of as the impact of infiltration. The same process is applied to the solar load through the windows and on the outside of the envelope. In the summer the solar load increases the energy consumption so subtracting the estimated impact of infiltration and both solar loads from the total yields an estimate of what energy conducts through the envelope insulation. These values are aggregated into a Home Efficiency Scorecard

where they also ranked against similar sized houses in the same climate zone. If infiltration is a large part of the total energy consumption that house would receive a low infiltration ranking. Advice and energy saving tips are provided for every ranking to give users ideas about how to improve their consumption and what would be the most efficient concern to address. The energy scorecard also shows how the HVAC power curve matches up against 14 and 21 SEER units. This analysis can also be performed for heating schedules during the winter season.

4.4.2 Additional Smart Meter Analytics

If only smart power meter data is available the HVAC consumption can be approximated using averages. The average percentage of total electricity usage that was attributed to AC (using the internet-connected thermostat and power curve derived from the smart meter data) for the 160 homes in Houston, TX with only one registered internet-connected thermostat in the summer of 2012 was 60%. However, Figure 4.24 shows how there is a large variance in this measurement.
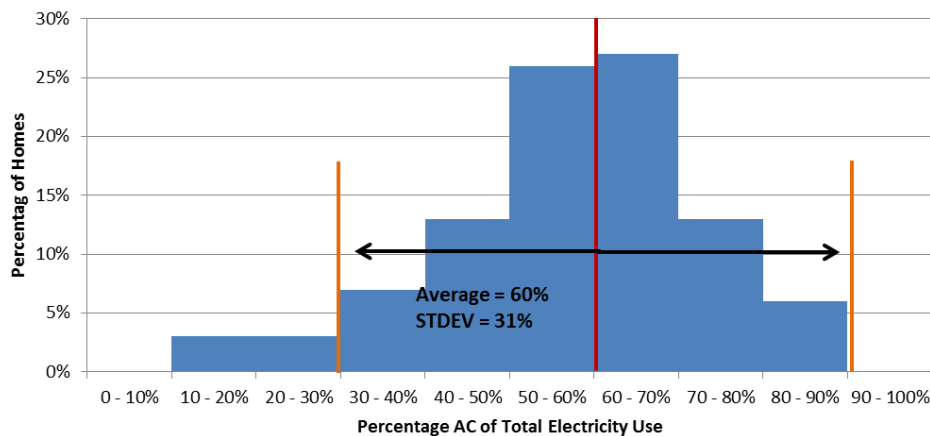
Figure 4.24: The percentage of total summer energy consumption attributed to AC for 250 houses in Houston, TX.

Figure 4.25 provides more insight into this by plotting the total energy consumed in the summer multiplied by 60% to approximate the AC usage versus the amount approximated using the thermostat method (the best substitution for actual AC usage). The Pearson product moment correlation coefficient for this set is 0.60 which suggests a weak correlation. The following method was designed to improve this uncertainty so more reliable recommendations can be made to homeowners about their AC usage.
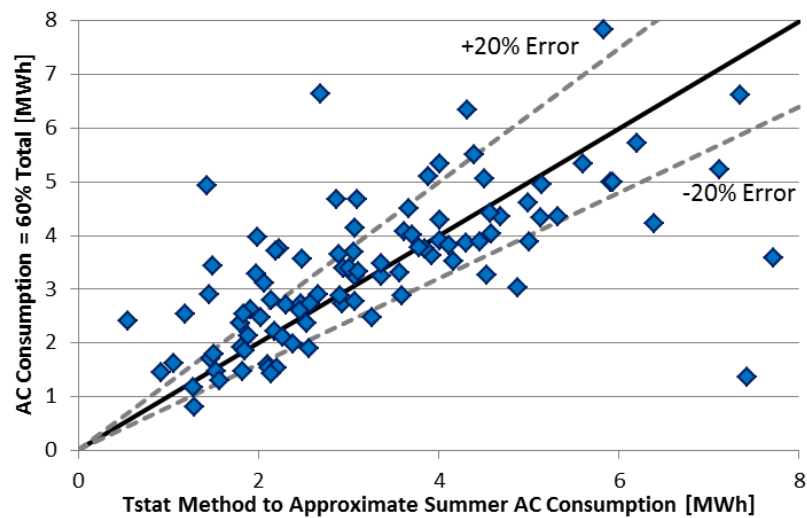
Figure 4.25: Comparing the 60% of the total load approximation for total seasonal AC consumption to what was computed using the thermostats.

This method uses a filtering technique to identify loads when internet-connected thermostat data are not available. Figure 4.26 plots the frequency of binned quarter hourly power meter measurements from the summer of 2012 for four houses from Houston, TX.
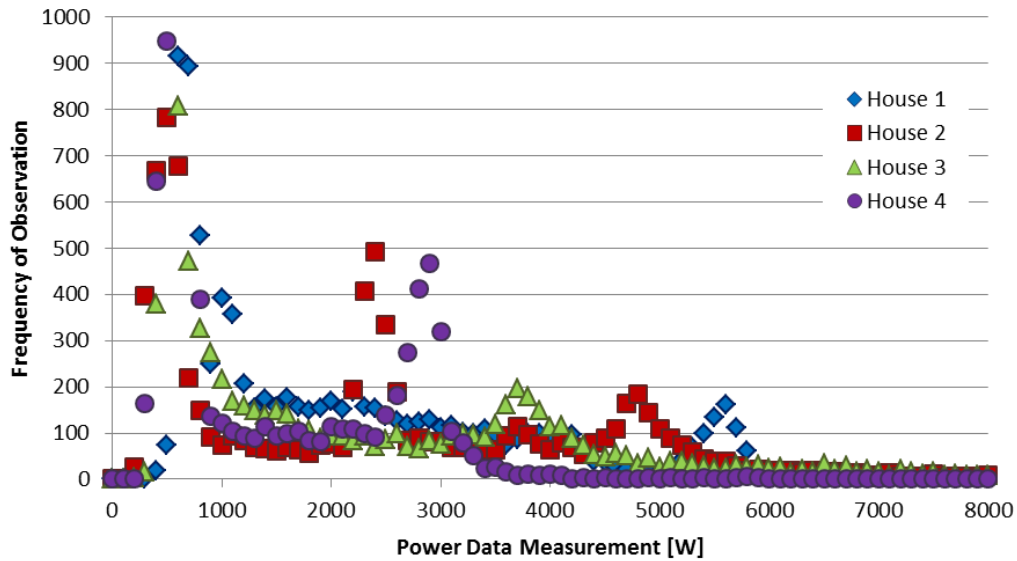
Figure 4.26: Frequency of observed quarter hourly power measurements of four houses from Houston, TX.

In Figure 4.26 distinct peaks in frequency can be seen at a range of a few power measurements below 1kW and at least one above 2.5kW. Air conditioners from these houses are expected to operate using more than 1.5kW, and account for a significant portion of the total load. This technique assumes that the bands below 1kW represent the base loads when no AC was ON and the band above 2.5kW was when the AC was ON. Both of these significant bands cover a range of power bins because of noise introduced from other non-base appliances and inconsistent quarter hourly AC duty cycles. A better estimate of the power bands of the base and AC loads is performed by filtering the data using the heat transfer metric. During low heat transfer the AC is not expected to cycle that frequently and the power band with the highest frequency would be the base load. When high heat transfer is observed the AC will be cycling more frequently and would be the highest frequency band that occurs. Low heat

117

transfer was defined as being observations below than the average minus half the standard deviation and high heat transfer was any observation larger than half the standard deviation above the average. Figure 4.27 plots the frequency of power measurements using the high and low heat transfer filter, and distinct bands around 1.7kW and 4kW are observed.
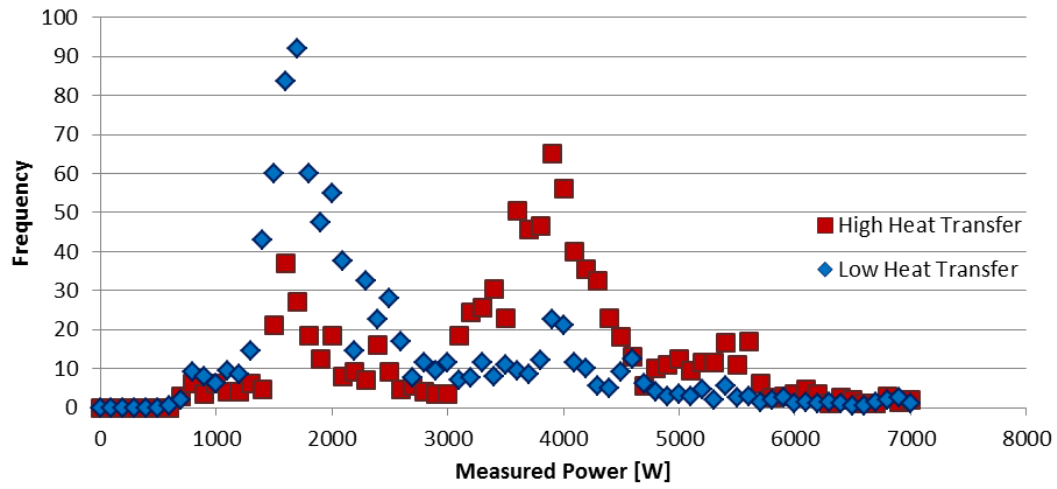
Figure 4.27: Frequency plot filtered using low and high heat transfer.

Figure 4.28 and Figure 4.29 aggregate the frequency plots of 50 houses (y axis) into a heat map where lower frequency power measurements are green and higher frequencies are red.
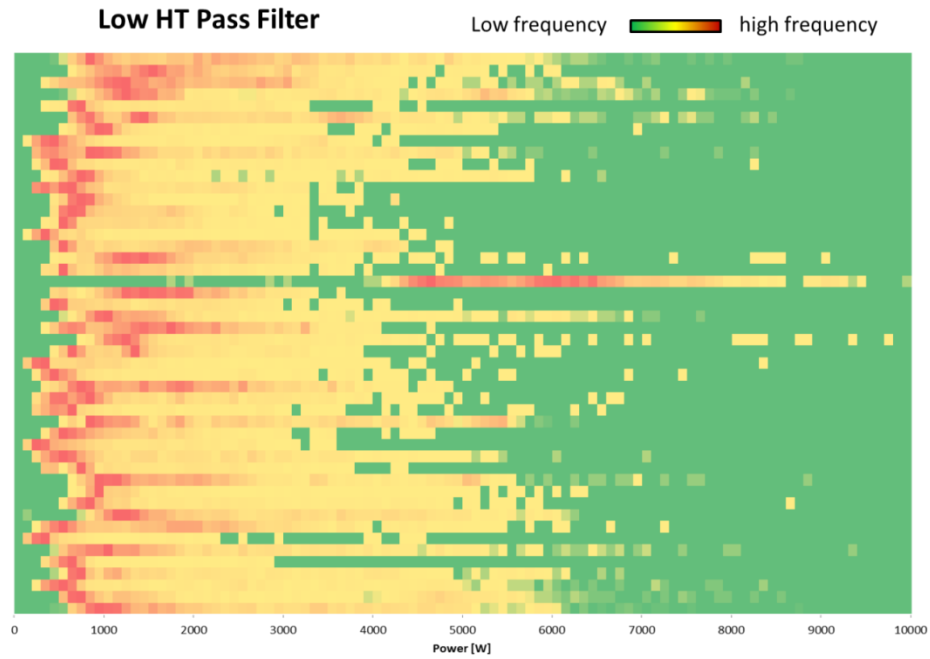
Figure 4.28: Heat map of the frequency of power measurements after filtering for low heat transfer for 50 houses in Houston, TX.
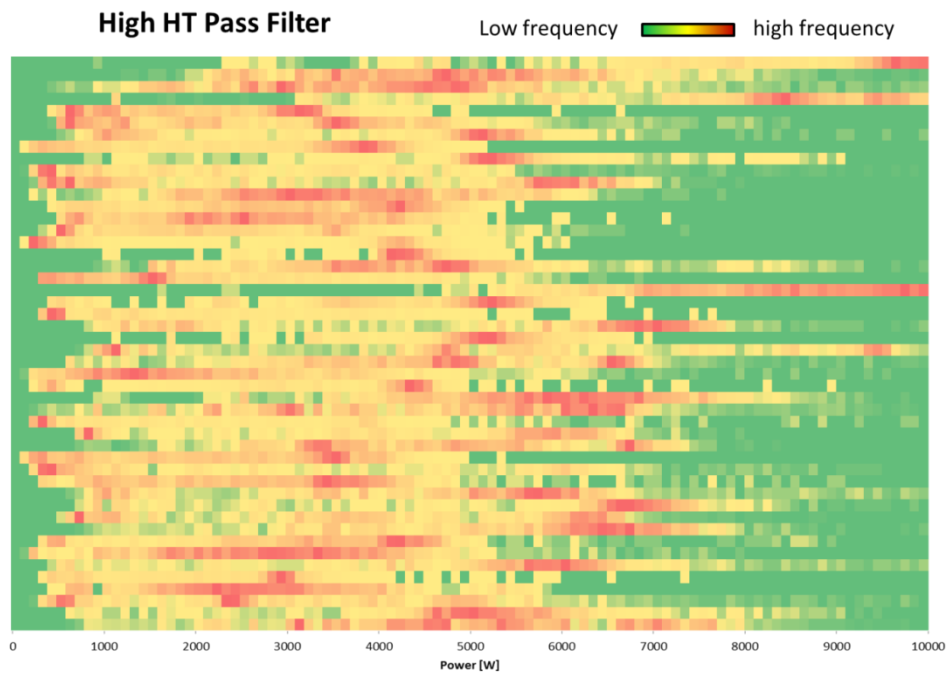
Figure 4.29: Heat map of the frequency of power measurements after filtering for high heat transfer for 50 houses in Houston, TX.

Once the approximate signature for each of these two loads is identified, the data is further filtered to disaggregate each load and an additional non-periodic appliance load. The non-periodic appliance loads are from devices like televisions, dryers, and ovens that do not occur regularly enough to cause a spike in the frequency like the base and AC loads. Anytime the load is larger than the base, but not large enough to be the AC, it is classified as a non-periodic appliance load. The same goes for any additional load above an expected AC load. Figure 4.30 depicts this process using a flow diagram.
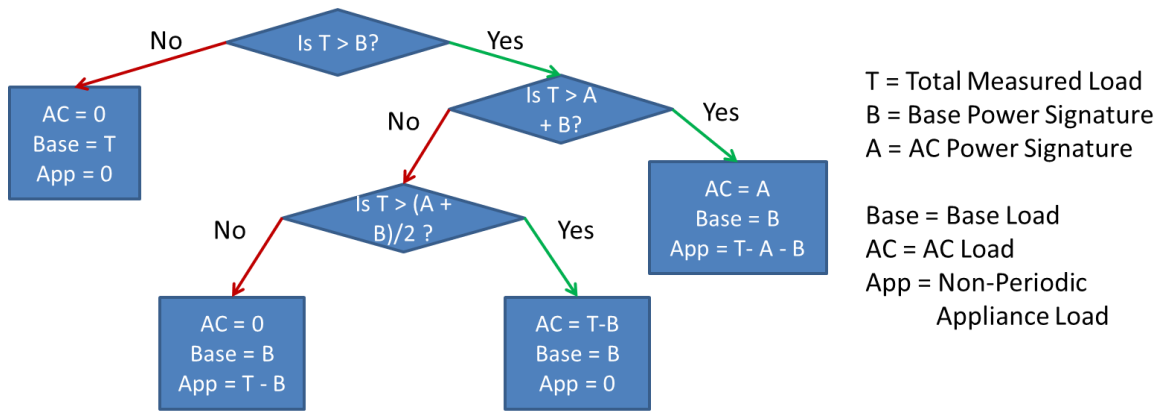
Figure 4.30: Flow diagram for disaggregating the total electricity load into base, non-periodic appliance, and AC loads.

Figure 4.31 and Figure 4.32 plot the hourly disaggregated energy consumption of houses in Houston, TX using this technique. Figure 4.33 summarizes the disaggregated energy consumption of the entire summer data of 50 houses from Houston, TX.
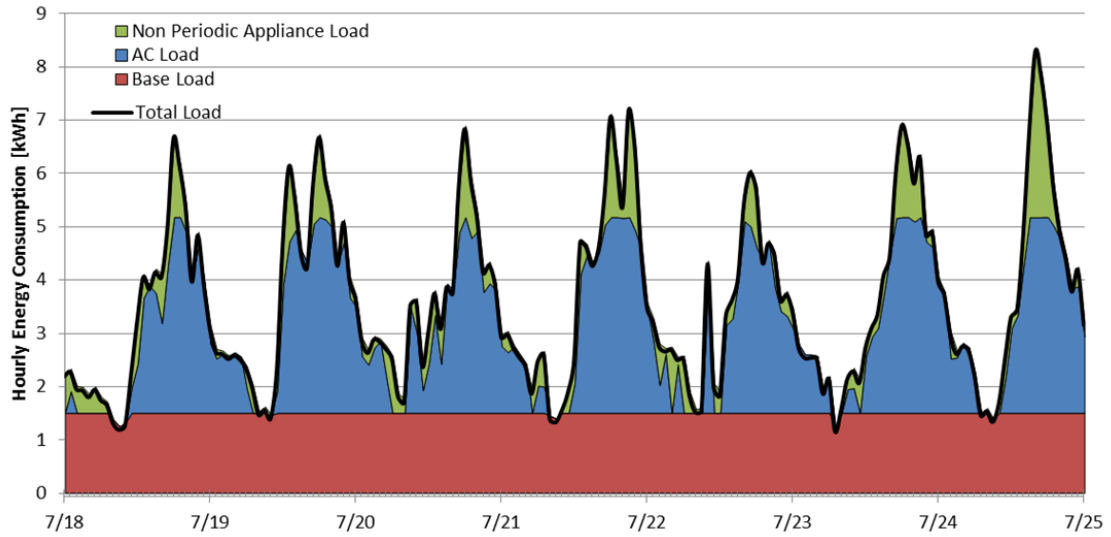
Figure 4.31: Disaggregated energy consumption for a house in Houston, TX using the heat transfer filtering method from July 18 to July 25, 2012.
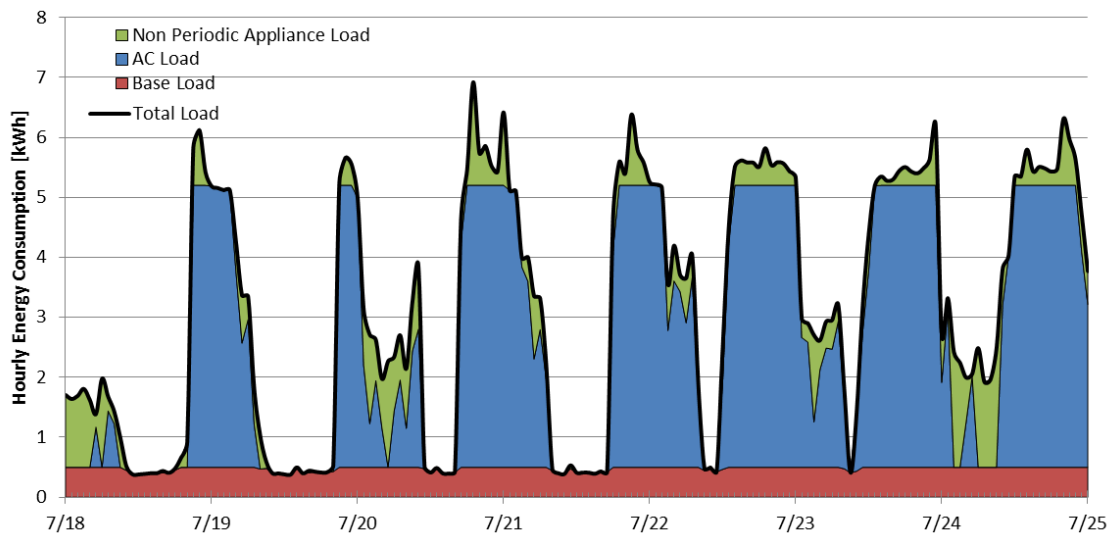
Figure 4.32: Disaggregated energy consumption for a second house in Houston, TX using the heat transfer filtering method from July 18 to July 25, 2012.
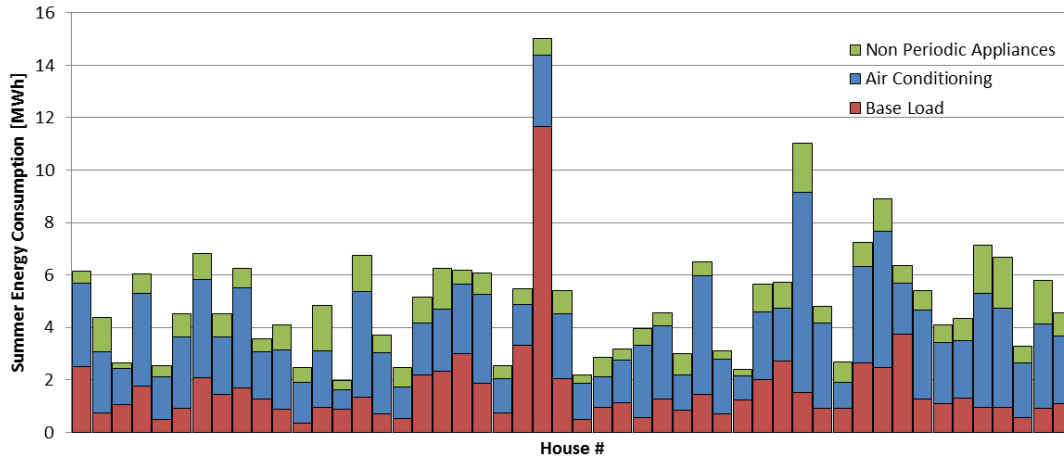
Figure 4.33: Disaggregated energy consumption for 50 houses from Houston, TX using the heat transfer filtering method from July 1 to September 30, 2012.

The average base and non-period appliance energy consumption of the 160 houses in Houston, TX were calculated to be 22% and 16% of the total using this method. The average AC consumption was 62% of the total. This was very close to the 60% average obtained when using the internet-connected thermostat data. Figure 4.34 shows an example day where both the filter and thermostat methods predict similar AC loads. However, not every user in the sample installed and registered internet-connected thermostats to control every air conditioner in their home. Figure 4.35 illustrates an example where there appears to be an additional AC cooling the home. The thermostat method only knows about the first AC installed and when it is turning ON. The filter method picks up the second AC because it cycles more frequently. This use case was not observed in the data that commonly, but could be the source of why the filter method calculates a higher AC percentage of the total load than the thermostat method.
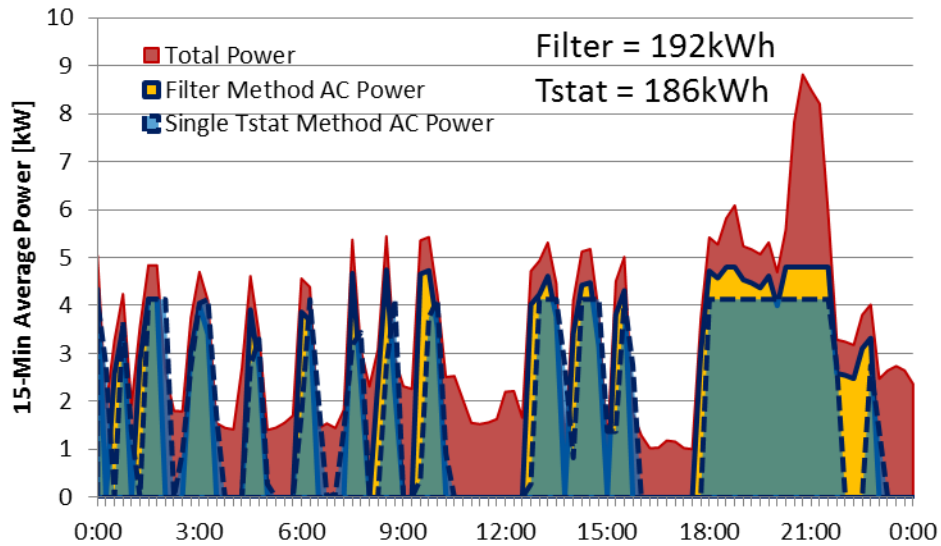
Figure 4.34: Disaggregation of AC load using the heat transfer filter and thermostat methods for a house in Houston, TX on August 22, 2012.
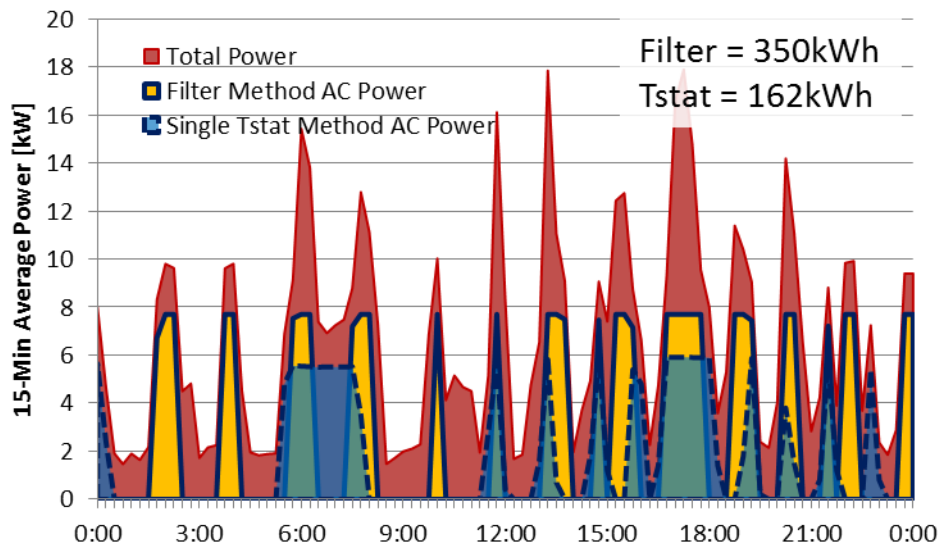
Figure 4.35: Disaggregation of AC load using the heat transfer filter and thermostat methods for a house in Houston, TX on August 3, 2012 that appears to have only registered a single thermostat.

Figure 4.36 plots the summer energy consumption of the houses in Houston, TX

using the filter method versus the thermostat method. The Pearson product moment correlation coefficient is 0.93 compared to the 60% total load estimate correlation coefficient of 0.60. The four outliers with larger filter method loads were investigated and three of them displayed data suggesting a second non-registered thermostat. Removing these from the dataset improved the correlation coefficient to 0.95.
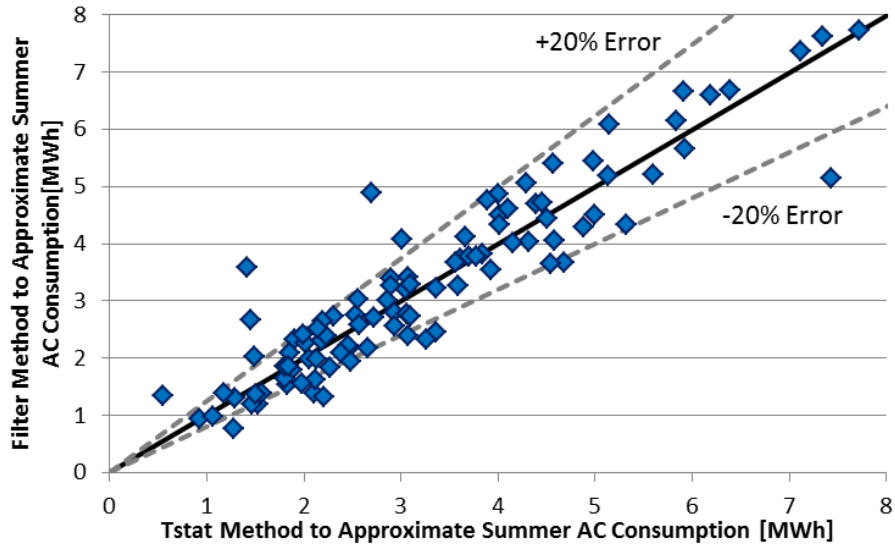
Figure 4.36: The approximated total energy consumed by the AC from July 1 to September 30, 2013 using the heat transfer filtering method versus the thermostat method.

Additional work needs to be performed to improve and test the validity of this method. Mathieu and Callaway (2012) present a method using Kalman filters that would make the decisions to disaggregate the load based on statistics and not just an arbitrary distance between the base and AC signature like the method presented. The method also needs to be adjusted to handle houses with multiple thermostats. It was observed, but not tested, that this method could be applied to electrical heating sources (heat pumps and electrical baseboards) during the winter.

Chapter 5: Conclusions

5.1 Thermostat Building Energy Model

This dissertation presented two residential building energy models and their applications for the energy market and homeowners. The thermostat building energy model used data from internet connected thermostats, local weather stations, and smart power meters to determine effective building parameters to represent the thermodynamics of a house. Genetic Algorithm was used to determine the building parameters that minimized the error between the predicted and actual indoor air temperature, but it was concluded that other solutions techniques may be more fitting and result in better performance. This model was designed to accurately forecast individual and aggregated loads quickly and require a minimal amount of stored data. The computational effort of the training phase was reduced by a factor of three by evaluating the 10 most extreme days (spider days) in a month set. It was observed that the building model aspects of daily load forecasting were executed on the order of milliseconds using a standard desktop computer.

This model was tested using data from 559 thermostats during heating (January 2013), cooling (July 2012), and a shoulder season (May-June 2012). The houses tested were from 83 different zip codes and five of the seven ASHRAE climate zones. In the 559 thermostat set the model averaged monthly ON time predictions only 1.9% higher than the measured data and an indoor temperature RMS error of 0.44°C. When combined with AC power curve modeling using smart meter data, the model predicted the daily RMS total energy of 250 additional houses in

Houston, TX from July 1 to September 30, 2012 3.4% better than the highest performing traditional prediction technique at an RMS error of 86 MWh. The hourly RMS error of 163 MWh was 28.2% better as well. The Pearson product-moment correlation coefficient between these predictions and the actual daily data increased from 0.911 for the outdoor temperature correlation prediction to 0.942. The hourly correlation coefficient increased from 0.883 to 0.931. This performance was based off of actual weather data, but when weather forecasts were used the model improved the daily prediction by 12.7% and hourly prediction by 5.4% with RMS errors of 89 and 191 MWh. The correlation coefficients improved over the best traditional method from 0.837 to 0.939 daily and 0.0872 to 0.913 hourly when using weather forecasts. The only statistically significant factor observed that influenced a model's ability to forecast was the thermostat setpoint schedule, and this is easily adjusted. The range in locations, climate, and HVAC operation successfully modeled in this study suggests that this model can be applied to any house nationwide.

5.2 Smart Meter Only Building Energy Model

The smart meter only building energy model was developed using smart power meter and local weather station data. This model was the most accurate at predicting general loads. An effective heat transfer from the outdoor to the indoor environment was correlated to past energy consumption to make predictions. Building parameters that minimized the error between the actual and predicted energy consumption were determined using the GA function, but could also be improved with a more appropriate technique. For 250 homes throughout the summer of 2012 in Houston, TX this model improved on the best tested traditional technique by 6.7% for the daily

prediction and 39.6% for the hourly prediction with RMS errors of 83 and 137 MWh. The daily and hourly correlation coefficients improved to 0.966 and 0.952 as well. It also improved the hourly predictions by 8.4% when weather forecasts were used with an RMS error of 185 MWh.

5.3 Applications

These building energy models were applied to forecast demand response capacity, generate more efficient thermostat setpoint schedules, and perform remote energy audits. The thermostat building energy model averaged forecasting the load of nearly 200 thermostats to within 5.9% during demand response events. This load encompassed the 2 hours before, during, and after three demand response tests in the late summer of 2012. The model was used to increase the total removed load during a demand response test on September 7, 2012 as well.

The thermostat building energy model was used as a simulation engine to test and implement energy saving thermostat setpoint schedules. Three mechanisms were developed and tested to improve the efficiency. Smart setback allows homeowners to set a desired indoor air temperature schedule, not simply a setpoint schedule. The developed algorithm determines new times for initiating setpoints that will meet the desired temperature most efficiently. Energy savings is realized by introducing more aggressive setbacks in schedules that previously had no or moderate setbacks. Setpoint smoothing starts setback periods earlier if the building model forecasts the indoor air temperature will not increase 1°C above the current setpoint. This mechanism reduces the time spent controlling to indoor temperatures further from the thermal equilibrium with the outdoor environment. Precooling takes advantage of the

higher cooling efficiency of colder outdoor air temperatures. These mechanisms are explained using cooling examples but the same logic can be applied for heating as well as when variable electricity rates are present.

Two methods were developed to disaggregate HVAC load from the total house electricity load. The first method used internet-connected thermostat and local weather station outdoor air temperature data to approximate a power curve. From the power curve the HVAC load during any duration of when it was ON was estimated. The second method correlated energy consumption to the heat transfer metric. This metric was used as a filter to identify the power signatures of base and HVAC load in smart meter data, and those signatures were used to further disaggregated the load. Simulations with the thermostat building energy model were performed to isolate how much energy was consumed by each major energy flow to maintain indoor comfort. The weather data were adjusted to separate the energy consumption attributed to solar loading, infiltration, and general heat diffusion through the envelope. This was presented to customers in a home efficiency scorecard to help educate them about their energy consumption.

5.4 Future Work and Recommendations

Several aspects of this project need to be improved and developed for its full potential to be realized. Every model tested and used in simulations was assuming that each thermostat controlled an isolated zone; however, nearly one third of all the thermostats in the Texas dataset were controlling different parts of the same structure. This general assumption of isolation is not valid for this subset. The model needs to be developed and tested on these multi-zone systems. Solution may involve additional

coefficients in the model that correspond to the conduction and convection between connected zones.

Utility and zone data is currently being recorded in Knight and Chincoteague Halls on the University of Maryland's campus. The goal is to adapt the energy models to buildings with 10-25 zones in an effort to use the same tools developed for single-zone residential units.

The data for the energy efficiency schedules needs to be analyzed to determine how much energy was saved and if the tools convinced users to be more aggressive with their setbacks. The algorithm can be improved to add in more mechanisms for energy saving. An idea is to use Fanger's thermal comfort (1970) as a metric instead of asking users for temperature setpoints. Higher-end model internet-connected thermostats measure both the indoor air temperature and the relative humidity. This can be paired with the wall and mass temperatures calculated in the model as a mean radiant temperature to create an approximation of thermal comfort. Internet-connected fans can also be added to provide effective cooling at lower energy costs. Controlling to thermal comfort may allow a larger temperature range for the scheduling algorithm to optimize. Optimizing schedules for heat pumps also can yield significant savings in climates that drop below freezing. During cold conditions the majority of the heat from heat pump units is provided from inefficient resistive heaters, and preheating during conditions where the coefficient of performance is higher can save energy.

The Home Efficiency Scorecard offers an effective remote energy audit. This needs to be developed further and tested against professional audits that require site

visits and blower door leak tests. A local auditing company has agreed to test a few local houses in the system. If the scorecard and remote audit are valid a web system or application will be developed to give any user with internet-connected thermostats and /or smart meters the ability to perform an audit, much like the Green Button Data (Irwin 2012).

This can be paired with a simulation tool already developed named the "WhatIffer." This tool allows users to test any setpoint schedule over any weather period on a web GUI or application. Essentially they can test a "what if I setback X" scenario to determine potential energy savings. The tool also has the ability to load in HVAC equipment and building structure upgrades. This gives users the ability to approximate the effectiveness of an upgrade and can help predict the payback period. The speed of this model lends itself to these type of simulations because an application such as this would only keep users engaged if it provides results instantaneously. A watered-down version can even be available to run on the thermostat to display the costs/benefits right at the source. Figure 5.1is a screenshot of the current beta-version web WhatIffer GUI.
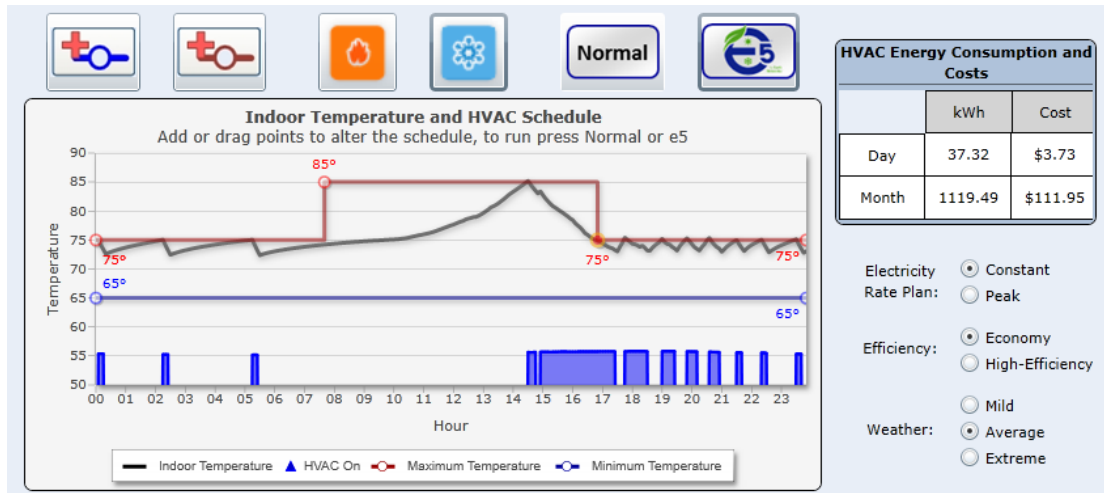
Figure 5.1: Current web WhatIffer GUI.

When implemented to predict forecasted and available loads, the models can be solved for the values that minimize the risk of a particular player in the energy market instead of just the forecasting ability. The costs associated with an error in the prediction may not be linearly related to the errors themselves. Under predicting on the average day may not be a problem for the market player, but under predicting on a hot day may be extremely costly. The errors may not even be due to the model but the forecasted weather used as an input. Therefore, the an approach like the stochastic dynamic programming model for analyzing optimal load estimation presented by Gabriel et al. can be taken (2004). Studying this can also be partnered with an investigation into how many modeled houses are required to represent the entire residential load. This would be useful information when weighing the costs and benefits of installing these thermostats and smart meters.

These energy models can also be used with economic models to optimize how the ISO can set prices to reduce peak demand and levelize costs. In a class project, the utility rates and how a customer would react was simulated as a Mixed Integer Non-

131

Linear Program. The thermostat building energy models were used to model the total load of several thousand homes given particular setpoint schedules. Because of the thermodynamic and desired comfort differences of houses, the ISO was able to levalize loads by offering various rate plans. This worked by assuming individual customers would always chose the schedule best for them, and periods of high consumption were effectively staggered. It would be worthwhile to investigate this more thoroughly and see if a model can be developed that improves the position of both the ISO and residential customers.

The human behavior aspects of DR can be explored using these building energy models and datasets. Customers may make different decisions about their energy consumption and participating in DR events given different economic incentives and backgrounds. Saving 20% on your energy bill may not be important for customers when the energy prices are low and the economy is doing well, but if money is tight customers may pay more attention to their consumption and want to energy tools. This information would be very valuable to the players in the energy market who need to forecast loads well into the future.

Appendices
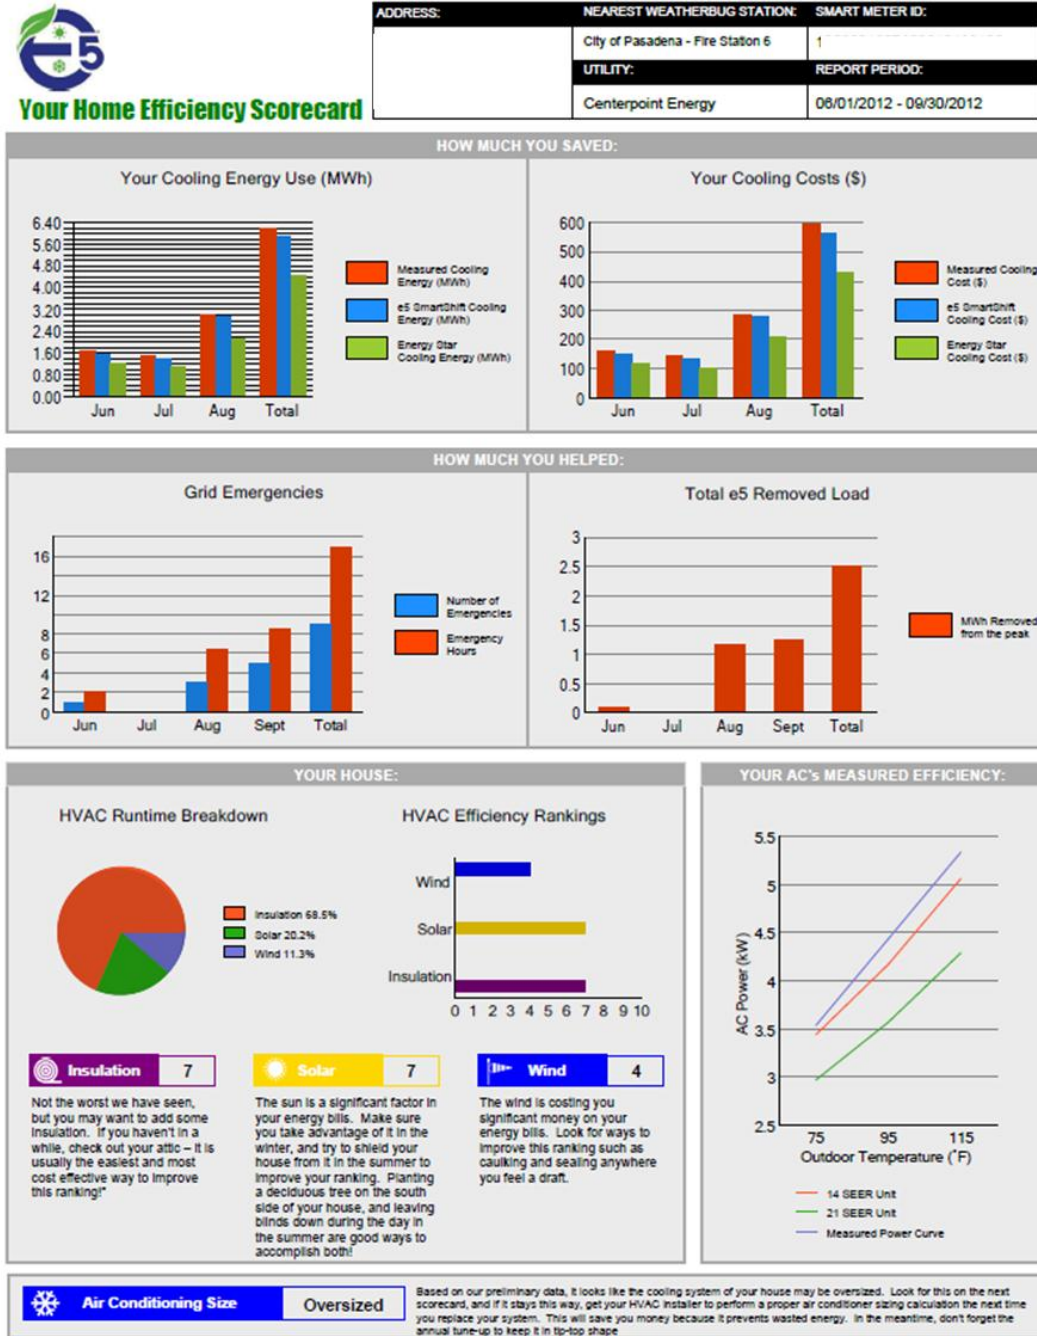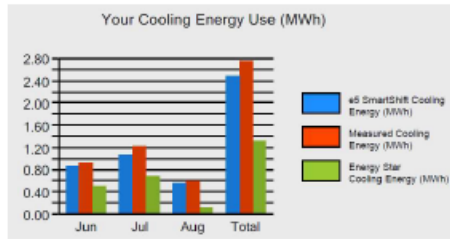
6.1 Example Home Efficiency Scorecard

Figure 6.1: Home efficiency scorecard example.

Your Home Efficiency Scorecard Explained

Your Cooling Energy Use:

Measured Cooling Energy - e5 uses your thermostat, smart meter and local weather station data to measure how much energy was consumed to run your air conditioning system. This is not based on averages and standards, but on your actual data and cutting-edge thermodynamic models to get a precise estimate of how much energy was used.

e5 SmartShift Cooling Energy - e5 develops a thermal model of your house and how it responds to varying weather conditions. Although the e5 thermostat optimization feature (SmartShift) was not available until August, we used your house model to calculate how much energy would have been used had e5 SmartShift been enabled throughout the summer.

Energy Star Cooling Energy - e5 can also use your house model to determine how much energy your air conditioning system would have used if you programmed your thermostat programmed according to the Energy Star recommended set points.
Energy Star Proper Use Guidelines for Programmable Thermostats
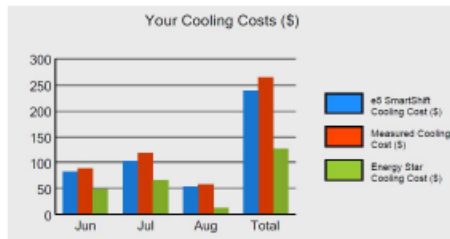
Your Cooling Costs:

e5 approximates your cooling energy costs based on your cooling energy use and the average price of electricity in you region ($0.097 per kWh). You can compare the "**Measured Cooling Cost**" in the chart to the "**e5 SmartShift Cooling Cost**" and "**Energy Star Cooling Cost**" to see how much money you can save. The costs shown correspond to a single thermostat's contribution. If you have multiple thermostats registered for e5, each will receive a separate report. If you have other non-e5 thermostats, or a window air conditioner, these costs may not represent your total cooling costs.

If a month is missing from either chart, then we did not have a sufficient amount of data for that period. This is typical for the first month in which you signed up for e5, or if there were connection issues with your thermostat.
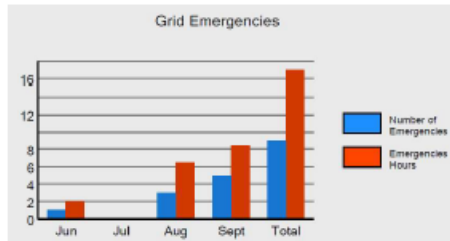
Grid emergencies:

There were several events during the summer known as "grid emergencies" where the electric grid was strained or we simulated a grid emergency for e5 testing. We shifted the air conditioning power used during these times by pre-cooling your house prior to the event, and raising your set point slightly during the event. This chart shows how many Grid Emergencies were initiated as part of the entire CenterPoint Summer Pilot Program and the duration each one. Your house participated in a subset of these events.
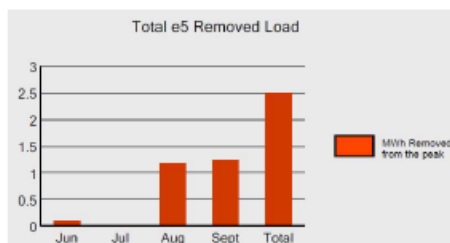
Total e5 Removed Load:

During the Grid Emergencies, energy was removed from the grid by reducing your air conditioning system consumption. This chart is not an individual thermostat breakdown, but rather a summary of the total energy removed by e5 during Grid Emergencies for all houses in the CenterPoint Energy e5 Summer Pilot Program, Your contribution went a long way in preventing blackouts, brownouts and higher rates. e5's main goal is to remove this load while still keeping you comfortable. If you did not notice these events or considered them tolerable, then we did our job. If you did notice them and/or you were negatively impacted, tell us about it at e5support@earthnetworks.com. We value your feedback as we work to expand the program in the future and offer customers more tools to ensure comfort.

Figure 6.2: Home efficiency scorecard explained P1

134

Your Home Efficiency Scorecard Explained

HVAC Runtime Breakdown
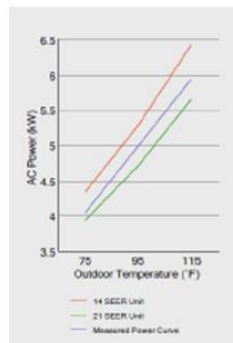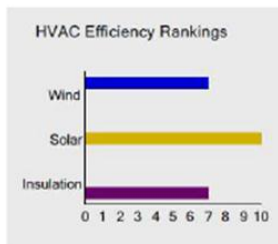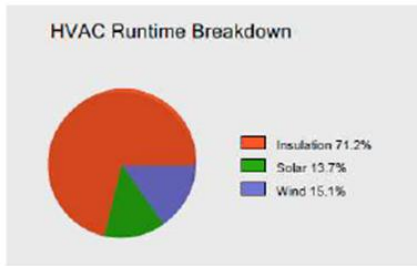
HVAC Runtime Breakdown:

e5 uses your house model to determine how much of your air conditioner runtime was due to each of the elements—temperature, sun and wind. These percentages give you a breakdown of where you cooling costs are coming from, and can help guide home improvements to reduce your energy consumption in the future.

The **Insulation Percentage** reflects how much energy was required to maintain the temperature difference between the inside and outside of your home over the course of the summer. A large Insulation Percentage is not necessarily a bad thing, and often means your other defenses are efficient. In hot climates like Texas, this will usually be the largest source of energy loss.

The **Solar Percentage** corresponds to the portion of cooling energy used to overcome the heat added to your house through the windows and walls by the sun. If this percentage is low, either your house is very efficient, or it is shaded. If it is not shaded, you can subtract this percentage from your "Measured Cooling Energy" total to estimate how much energy you would use if it was always cloudy or your house was well shaded.

The **Wind Percentage** shows the cooling energy required to cool the hotter outdoor air that infiltrated into your house through cracks and leaks. A big portion of this can be due to open windows and doors (in which case now you know how much you are paying to cool the neighborhood instead of just your house!).

HVAC Efficiency Rankings:

e5 uses your house model to determine how efficient your house is at repelling the elements compared to other similarly sized houses in the same climate region. Note that these rankings are not dependent on your thermostat set points, but a reflection of how efficient your house is compared to other similar e5 participants. The rankings range from 0 for inefficient to 10 for a very efficient house, with value of 5 being the average house.

The **Insulation Ranking** ranks the efficiency of the wall and attic insulation at preventing outside heat from entering your house. A low ranking here suggests insufficient levels of insulation and/or degradation of your building's envelope.

The **Solar Ranking** ranks the efficiency of the house at stopping solar radiation from heating your indoor environment. A low ranking here suggests a large portion of your house is subject to direct sunlight and/or your window blinds are not sufficient.

The **Wind Ranking** ranks the efficiency of the house at preventing hotter outdoor air from infiltrating inside. Summer infiltration is less of a concern than in winter, but a poor ranking here could result from leaving windows and doors open.

Your AC's measured efficiency:

The efficiency of an air conditioner is determined by how much cooling energy it produces given the amount of electrical energy input. Air conditioning manufacturers use a calculated Seasonal Energy Efficiency Ratio (SEER) to rate units. Higher SEER systems are more efficient and require less energy and money to operate. As the outside air temperature rises during the day, air conditioners consume more power to provide the same amount of cooling. A power curve is determined by measuring how much power the air conditioner requires at various outdoor air temperatures. e5 generates your **Measured Power Curve** using your thermostat, smart meter and closest weather station data. The plot shows how your Measured Power Curve compares against new efficient (21 SEER) and less efficient (14 SEER) units estimated to be the same capacity as yours. The higher and steeper the curve, the less efficient your air conditioner is.

Your AC's measured efficiency:

e5 uses your house model and thermostat data to simulate the energy consumed by your air conditioner for mild, seasonal average, and extreme weather conditions. This ranking estimates how well your air conditioning system is sized to your house. "Oversized" systems can keep you comfortable during the most extreme heat but may waste energy during normal operation. An "Undersized" ranking may not be a bad thing if you can live with the level of comfort it provides, but a low ranking could also be triggered by degradation of your system. This naturally occurs over time and may be worth a checkup, especially if your system is over 8 years old.

Figure 6.3: Home efficiency scorecard explained P2

An updated version of the home efficiency scorecard that includes further load disaggregation and bill forecasting is shown in Figure 6.4.

Figure 6.4: Home efficiency scorecard update April 2013.

6.2 Additional Plots from Section 3.2

The following section contains plots from the study performed in section 3.2. Figure 6.5 plots the average hourly HVAC duty cycles and indoor air temperatures observed in the data and predicted for the entire set of models from May 15 to May 22, 2012.

Figure 6.5: Hourly average HVAC duty cycles and indoor temperatures from May 15 to May 22, 2012.

Figure 6.6 plots the average duty cycle and indoor temperature for every 5-minute interval in the May-June dataset.

Figure 6.6: Average HVAC duty cycle and indoor temperature for every 5-minute interval in the May-June dataset.

Figure 6.7 plots the average hourly HVAC duty cycles and indoor air temperatures observed in the data and predicted for the entire set of models from July 1 to July 8, 2012.

Figure 6.7: Hourly average HVAC duty cycles and indoor temperatures from July 1 to July 8, 2012.

Figure 6.8 plots the average duty cycle and indoor temperature for every 5-minute interval in the July dataset.

Figure 6.8: Average HVAC duty cycle and indoor temperature for every 5-minute interval in the July dataset.

Figure 6.9 plots the average hourly HVAC duty cycles and indoor air temperatures observed in the data and predicted for the entire set of models from January 1 to January 8, 2013.
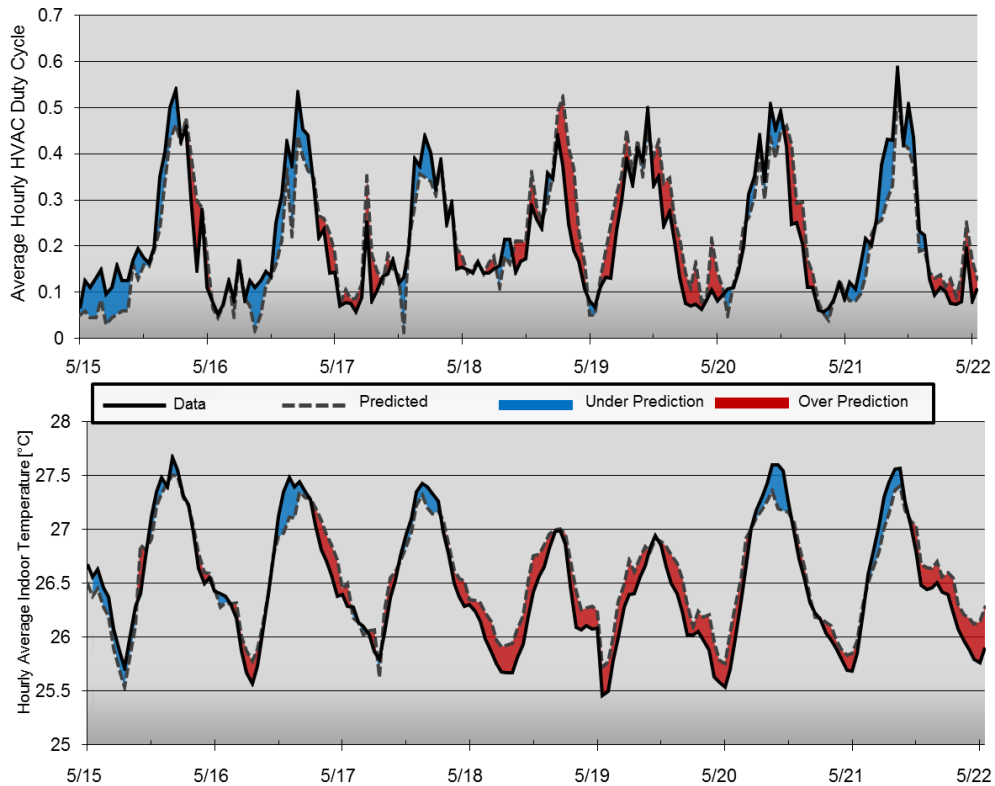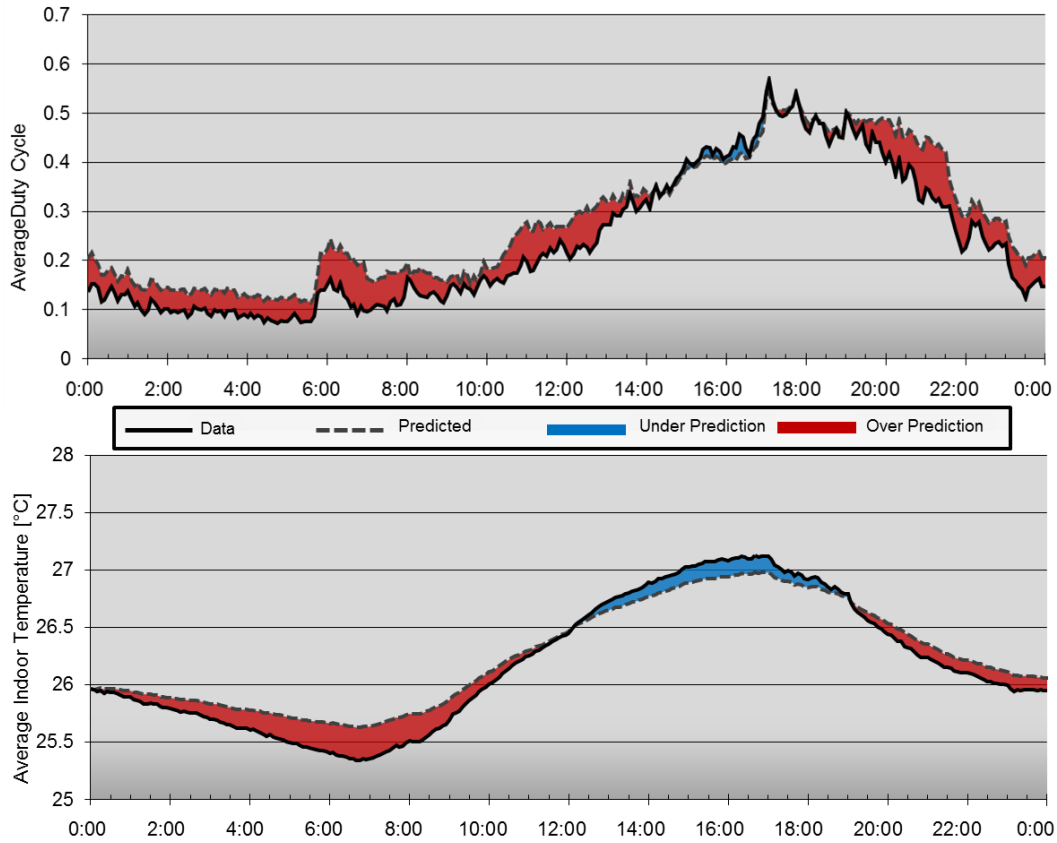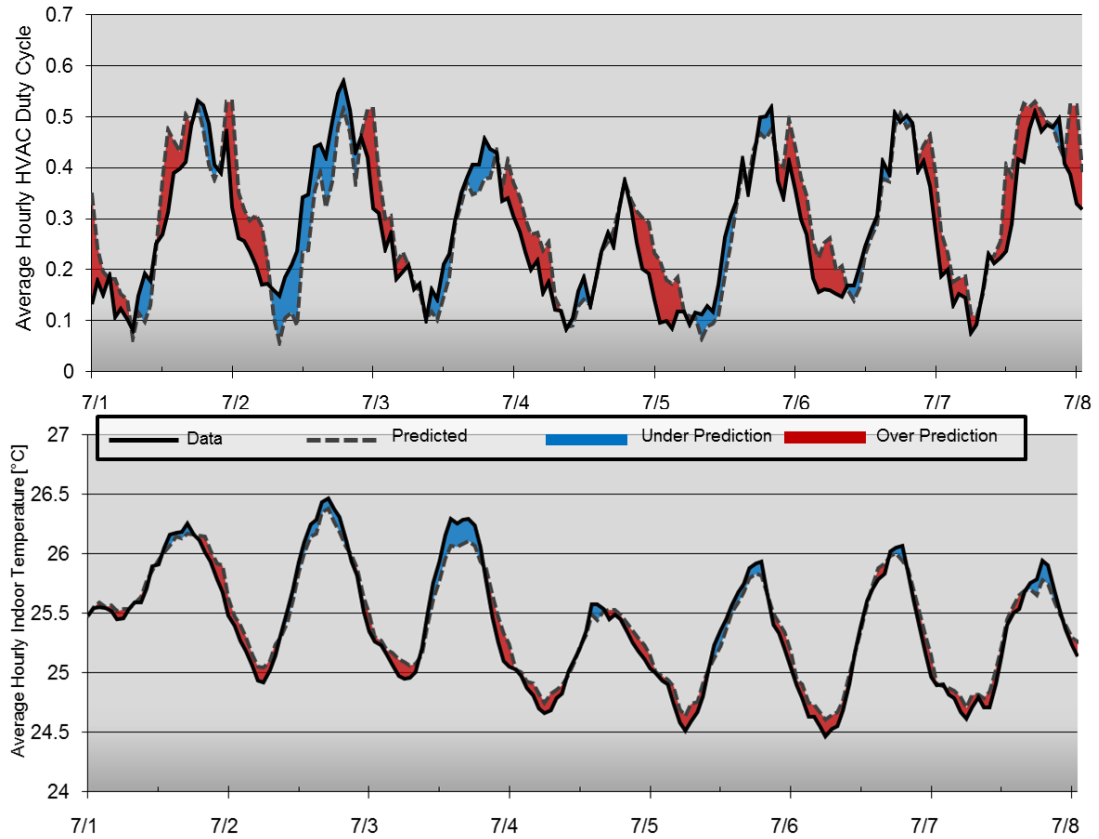
Figure 6.9: Hourly average HVAC duty cycles and indoor temperatures from January 1 to January 8, 2013.

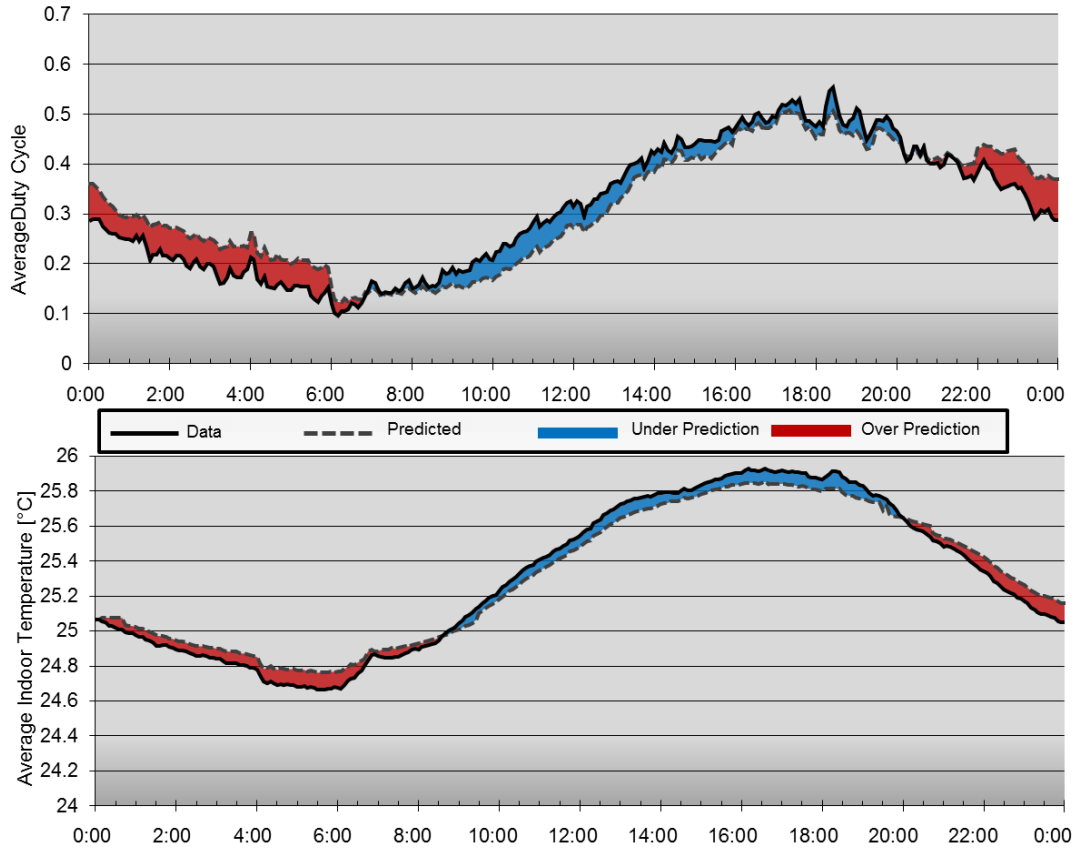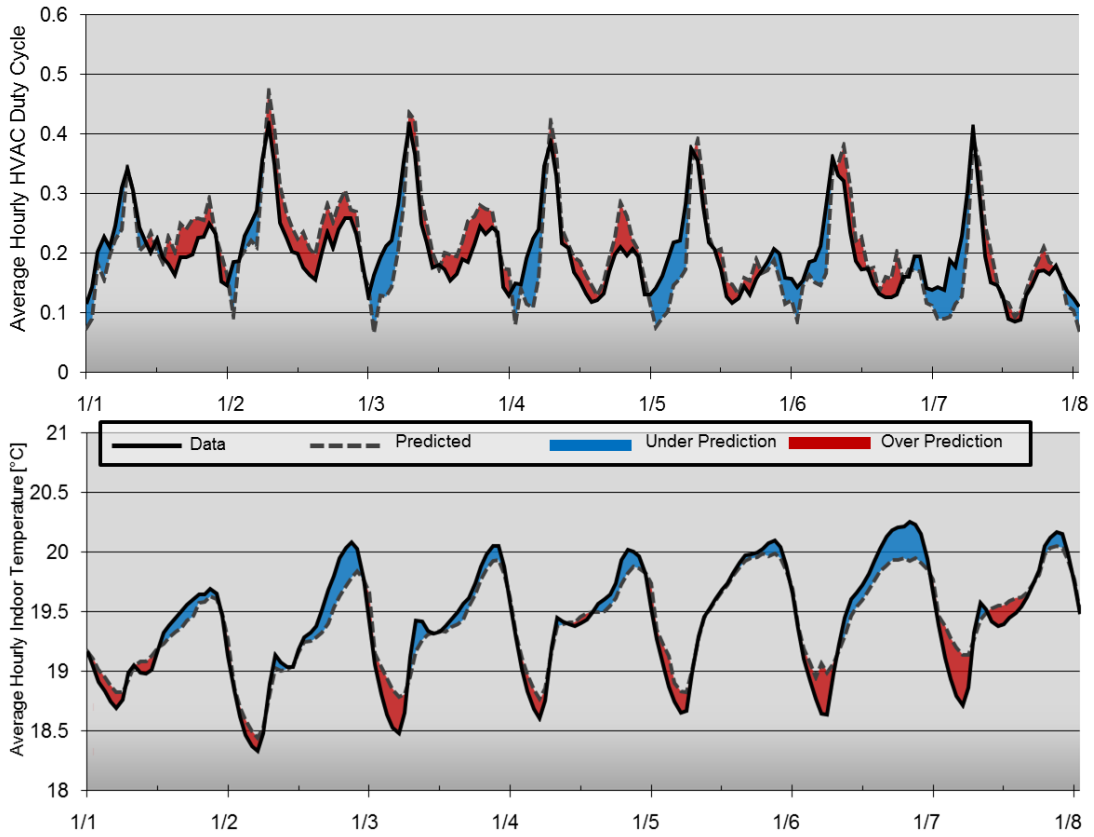Figure 6.10 plots the average duty cycle and indoor temperature for every 5-minute interval in the January dataset.
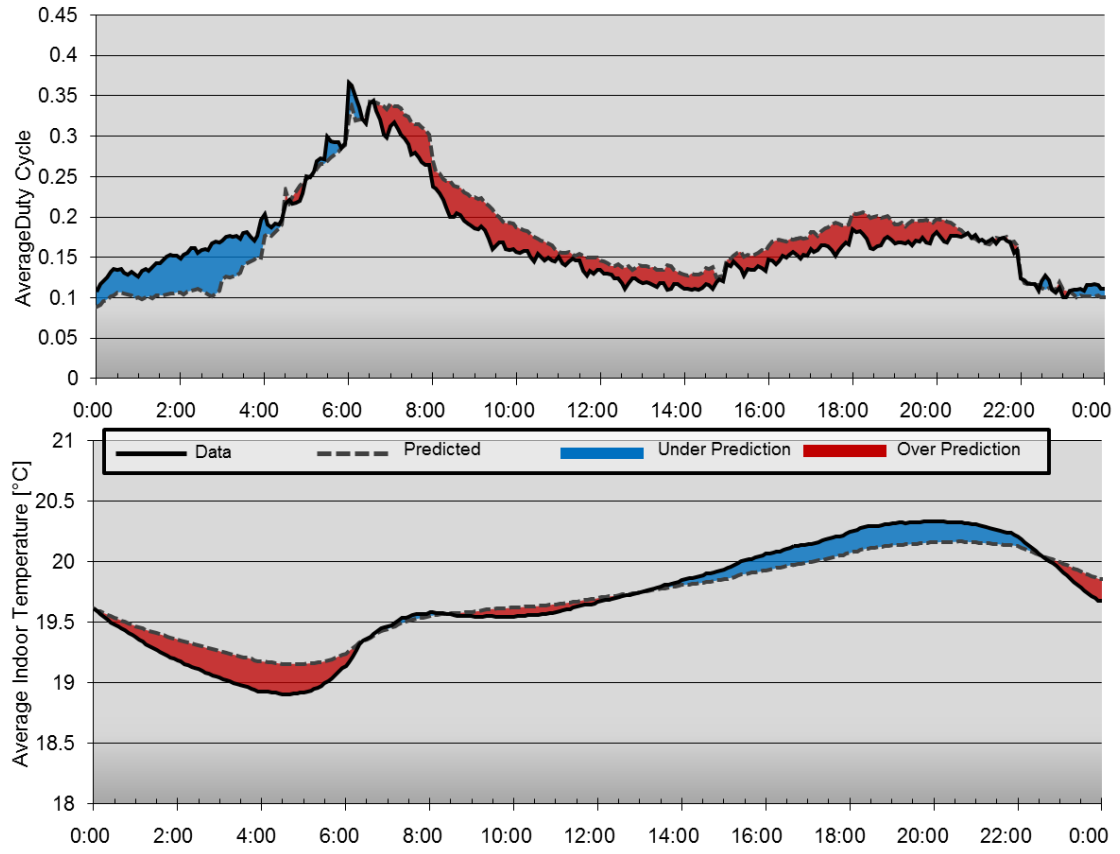
Figure 6.10: Average HVAC duty cycle and indoor temperature for every 5-minute interval in the January dataset.

Glossary

Complete ON/OFF Switch - Instances when the HVAC goes from being fully OFF for the defined interval to fully ON for the interval or vice versa.

Control Phase – Simulation where the thermostat building energy model is operated without the knowledge of when the HVAC system turned ON. In place of this a control algorithm using the setpoint and the last calculated indoor air temperature predicts the ON time.

Cooling Degree Day (CDD) - A measurement designed to reflect the demand for energy needed to cool a building on a given day. Traditionally, CDD are calculated using the high and low daily outdoor air temperature and a base of 15.5°C below which cooling would not be required.

Demand Response (DR) - Strategies that reduce the demand for power during peak times by shifting or curtailing its use.

E5 – Earth Networks Program for DR and Energy Efficiency < http://earthnetworks.com/e5.aspx>

Hysteresis – A temperature range above (cooling) or below (heating) the setpoint where the thermostat will not turn ON the HVAC system. Hysteresis is used to prevent the HVAC system from cycling too frequently to maintain an exact temperature. The thermostats used in this study defaulted to 0.56°C, but could be set in increments of 0.28°C.

Power Curve – A correlation of the power drawn by the compressor of an HVAC system to the condensing (AC) or evaporating (heat pump) outdoor air temperature. The three Carrier units referenced in this dissertation reported linear power curves in their data sheets.

Precool- To run the AC earlier than dictated by the thermostat to store cooling energy in the home. Precooling is used when preparing for demand response events, to reduce the usage during periods of high electricity prices, and when it is more efficient than normal operation. The same results can be obtained by preheating when applicable.

Recovery Phase – Period after a demand response event when thermostat setpoints are returned to the standard operation.

Root Mean Square (RMS) – Measure of the magnitude of a varying quantity. In this dissertation RMS error refers to the sum of the absolute value of the differences between a prediction and the actual measured value evaluated at every point of the defined range. This RMS error is both averaged and left as the total sum, and is used as performance metrics for many evaluations.

Setback – Raising (cooling season) or reducing (heating season) the thermostat setpoint temperature to be in closer thermodynamic equilibrium with the outdoor environment in an effort to reduce operational energy. This is often performed while the house is unoccupied.

Setpoint – The indoor air temperature setting the thermostat controls the HVAC system to maintain for a given period.

Smart Meter Only Building Energy Model – Grey-box building energy model that is trained using local weather and smart power meter data. The model correlates an effective heat transferred from the outdoor to indoor environment to power data.

Smart Power Meter – Electricity meter that records data in shorter intervals than traditional analog and observed metering. Smart meters communicate with the utility directly for monitoring and billing purposes.

Spider day – One of 10 of the most extreme days in a month data set that can be used to approximate the entire set.

Thermostat Building Energy Model – Grey-box building energy model that is trained using internet-connected thermostat, local weather, and smart meter power data. The thermodynamic equations are represented by 12 parameter coefficients.

Bibliography

"Amazon DynamoDB (beta)." Amazon DynamoDB. N.p., n.d. Web. 02 Apr. 2013. <http://aws.amazon.com/dynamodb/>.

ASHRAE 2006. ASHRAE Standard 169-2006 Weather Data for Building Design Standards. Atlanta, GA: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2009. ASHRAE Handbook-Fundamentals. Atlanta, GA: American Society of Heating Refrigerating and Air Conditioning Engineers, Inc.

Braun, J.E., Chaturvedi, N. 2002. An Inverse Gray-Box Model for Transient Building Load Prediction. HVAC&R Research. 8 (1).

"CONTROL." RADIO THERMOSTAT » Welcome. N.p., n.d. Web. 02 Apr. 2013. <http://www.radiothermostat.com/control.html>.

Degelman, L. 1991. "A statistically-based hourly weather data generator for driving energy simulation and equipment design software for buildings", Proc. Building Simulation '91, Int'l Building Performance Simulation Assoc. (IBPSA), 20-22 Aug., pp 592-599.

Degelman, L. 2004. "Simulation and uncertainty: weather predictions", Advanced Building Simulation,(Malkawi and Augenbroe, Ed.), Spon Press, New York and London, Chap. 3, pp 60-86.

Fanger, P. O. (1970). Thermal comfort. Analysis and applications in environmental engineering. Thermal comfort. Analysis and applications in environmental engineering.

Henze, G.P., Kalz. D.E., Felsmann, C., Knabe, G.. 2011. Impact of Forecasting Accuracy on Predictive Optimal Control of Active and Passive Building Thermal Storage Inventory. HVAC&R Research, 10(2).

Herter, K., Wayland, S. 2010. Residential response to critical-peak pricing of electricity: California evidence. Energ. 35: 1561-1567.

"Home Energy Reports." Opower. N.p., n.d. Web. 02 Apr. 2013. <http://opower.com/what-is-opower/reports>.

Incropera, Frank P., and David P. DeWitt. Fundamentals of Heat and Mass Transfer. New York: Wiley, 1990. Print.

Irwin, C. "Energy.gov." Energy.gov. N.p., n.d. Web. 18 Sept. 2012. <http://energy.gov/articles/green-button-data-more-power-you>.

Kay, L F., and J Baughman. "Consumer Advocates Call for Review of BGE Program." Baltimore Sun. N.p., 25 July 2011. Web. 18 Sept. 2012. <http://articles.baltimoresun.com/2011-07-25/news/bs-md-bge-peak-rewards-psc-20110725_1_mark-d-case-peakrewards-customers-cycling>.

Kempton, W., Reynolds, C., Fels, M., Hull, D. 1992. Utility control of residential cooling: resident-perceived effects and potential program improvements. Energy and Buildings. 18: 201-219.

Li, Q., Meng, Q., Cai, J., Yoshino, H., Mochida, A. 2009. Predicting hourly cooling load in the building: A comparison of support vectormachine and different artificial neural networks. Energy Conversion and Management. 50: 90–96.

Lowry, G, Man-Wai, L. 2004. Modelling the passive thermal response of a building using sparse BMS data. Applied Energy, 78: 53-62.

Mathieu, J.L.; Callaway, D.S., "State Estimation and Control of Heterogeneous Thermostatically Controlled Loads for Load Following," System Science (HICSS), 2012 45th Hawaii International Conference on , vol., no., pp.2002,2011, 4-7 Jan. 2012

Mathieu, Johanna L., Duncan S. Callaway, and Sila Kiliccote. "Variability in Automated Responses of Commercial Buildings and Industrial Facilities to Dynamic Electricity Prices." Energy and Buildings (2011): n. pag. Print.

MATLAB version 7.10.0. Natick, Massachusetts: The MathWorks Inc., 2010.

Neto, A.H., Fiorelli, F.A.S. 2008. Comparison between detailed model simulation and artificial neural network for forecasting building energy consumption. Energy and Buildings. 40: 2169–2176.

Newsham, G.R., Birt, B.J., Rowlands, I. H. 2011. A comparison of four methods to evaluate the effect of a utility residential air-conditioner load control program on peak electricity use. Energy Policy. 39: 6376-6389.

Newsham, G.R., Bowker, B.G. 2010. The effect of utility time-varying pricing and load control strategies on residential summer peak electricity use :A review. Energy Policy. 38: 3289-3296.

"NSRDB: 1991- 2005 Update: TMY3." NSRDB: 1991- 2005 Update: TMY3. N.p., n.d. Web. 04 Apr. 2013. <http://rredc.nrel.gov/solar/old_data/nsrdb/1991-2005/tmy3/>.

Oussar, Y., Dreyfus, D. 2001. How to be a graybox: dynamic semiphysical modeling. Neural Networks. 14: 161-172.

Patankar, Suhas V. Numerical Heat Transfer and Fluid Flow. Washington: Hemisphere Pub., 1980. Print.

Peng, C., Wu, Z. 2008. Thermoelectricity analogy method for computing the periodic heat transfer in external building envelopes. Applied Energy. 85: 735-754.

Pionergy Consulting, e5 Demand-Response Program Pilot Sampling Approach to Estimate Deemed Savings, September 11, 2012. Print.

"Smart Grid RECOVERY ACT SMART GRID PROGRAMS." SmartGrid.gov:. N.p., n.d. Web. 18 Sept. 2012. <http://www.smartgrid.gov/recovery_act/overview>.

"Smart Power." Business Money Smart Power Comments. N.p., n.d. Web. 02 Apr. 2013. <http://business.time.com/2013/03/28/smart-power/>.

Tan, K.C., Li, Y. 2002. Grey-Box model identification via evolutionary computing. Control Engineering Practice. 10: 673-684.

"Thermostat Guidelines." US Department of Energy: Energy Star. N.p., n.d. Web. 02 Apr. 2013. <http://www.energystar.gov/index.cfm?c=thermostats.pr_thermostats/>.

Walpole, Ronald E. Probability & Statistics for Engineers & Scientists. Upper Saddle River, NJ: Pearson Prentice Hall, 2007. Print.

"Weather Network." Earth Networks. N.p., n.d. Web. 18 Sept. 2012. <https://www.earthnetworks.com/OurNetworks/WeatherNetwork.aspx>.

Wilcox, S. and W. Marion. User's Manual for TMY3 Data Sets, NREL/TP-581-43156. April, 2008. Golden, Colorado: National Renewable Energy Laboratory.

Wu, S., Sun, J.Q. 2012. A physics-based linear parametric model of room temperature in office buildings. Building and Environment, 50: 1-9.

Zhao, H., Magoules, F. 2012. A review on the prediction of building energy consumption. Renewable and Sustainable Energy Reviews, 16: 3586-3592.