A CASE FOR THERMOSTAT USER MODELS

Bryan Urban and Carla Gomez
Fraunhofer Center for Sustainable Energy Systems, Cambridge, MA, USA
Email: burban@fraunhofer.org

ABSTRACT

Today's thermostat setpoint models naïvely assume fixed schedules, ignoring the reality of user control and its large variability. Better models must include more realistic user-behavior profiles to correctly evaluate the energy benefits of the next generation of thermostats against a realistic baseline. Data from a recent thermostat field study were analyzed to demonstrate the variation and patterns associated with manual adjustment of programmable thermostats and its consequences on observed and simulated energy consumption. A practical modeling technique for describing variable setpoint schedules was applied and compared with standard, fixed setpoint assumptions. Room air temperature data from 63 apartments near Boston, MA were used to generate unique hourly heating setpoint schedules. These observed temperature histories were then used to model the expected variation in energy use of an apartment due to manual thermostat adjustment. Significant differences in energy consumption were observed when variable setpoints were used instead of fixed setpoints, indicating the need for improving thermostat assumptions and updating models with more realistic schedules.

INTRODUCTION

The influence of occupant behavior on thermal energy consumption has been discussed for over 30 years, yet the standard practice when simulating thermal controls is to assume either a constant setpoint or a fixed thermostat schedule (Urban et al. 2012). When modeling programmable thermostats, the ASHRAE 90.2 standard requires modelers to use constant set-back and set-up temperatures of 60°F and 68°F for residential heating (ASHRAE 2007). This can be problematic. Not only are these and other typical setpoint temperature assumptions unrealistic, they fail to capture the frequent variability caused by humans pushing buttons.

We already know that users often misinterpret how their thermostat works (Kempton 1986; Nevius & Pigg 2000; Meier et al. 2010; Peffer et al. 2011; Malnick et al. 2012) and perhaps 20% of programmable thermostat users fail to set the correct time (Meier et al. 2011). Common misconceptions about programmable thermostats include:

- Thermostat is an on/off switch
- Thermostat is a dimmer switch
- Thermostat is an accelerator
- Turning down the thermostat has little or no effect on energy consumption

Such misconceptions and misuse can lead to unexpected usage patterns.

Operationally, users normally have several options when interacting with the thermostat including a program, a temporary override, and a long-term override or permanent hold. A program is an automated schedule of setpoints that vary in time. Thermostats commonly have an energy saving program as a default, and users may manually specify their own program. A temporary hold is a manual adjustment that overrides the program temporarily, typically for a few hours or until the next schedule period arrives, when the program resumes. A long-term hold is a manual setpoint adjustment that overrides the program permanently until it is manually disabled. Beyond these standard options, thermostats may also be used to switch the heating system on or off entirely. Newer thermostats include features that allow for more detailed programs, web-enabled control, and automatic setback functions.

Thermostats primarily aim to help occupants remain thermally comfortable. Secondly, thermostats may enable reduced HVAC energy consumption, e.g. through the use of temperature setbacks. Whether or not energy savings are realized depends entirely on how a thermostat is used. This fact has led to some confusion and disagreement about the energy saving potential of thermostats. Occupant behavior, in response to varied environmental conditions and occupancy patterns, can produce energy savings or lead to excess consumption. Even the mere presence of a programming feature may induce people to use more energy (Malnick et al. 2012). Better simulation in conjunction with more field studies will help resolve some of this confusion.

One reason modelers may be tempted to ignore thermostat usage variability is that it adds complexity to the models. This need not be the case. In 2012 we presented a straightforward method for applying variable thermostat schedules in simulations using popular building energy simulation tools (Urban et al. 2012), and other approaches for doing so are referenced therein. The real barrier exists in knowing how typical occupants actually use their thermostats. The models exist; the appropriate assumptions are not well understood. Relatively few studies have examined how thermostats are really used in homes, and results are inconsistent, so modelers have little feedback about how realistic their current approaches really are.

In this first part of this study we will investigate the observations from a field study to expose several categories of user behavior types. In the second part, we will simulate the observed temperature schedules from the field study using a prototypical apartment building. This series of simulations will illustrate the isolated effect behavior may have on simulated heating energy consumption.

A Thermostat Field Study

We recently completed a study to gain a better understanding of how occupants use thermostats (Sachs et al. 2012). In the winter of 2011, we installed programmable thermostats in 82 residential units in a masonry structured apartment building in Revere, MA. Additionally, we installed data loggers in these units to collect room temperature at 10 minute intervals and HVAC state changes as they occur. We also recorded the heating energy consumption from the gas meter for each unit on a weekly basis. Of the 82 units tested, 60 units provided adequate data for this original study. Based on data gathered during the winter of 2011-12, we will show what real-world thermostat operation looks like for a range of occupants in a single building.

Since manual thermostat adjustment is only one component of heating energy use, we must remain aware that, in addition to behavior, other building parameters influence thermal gains and losses, and most prominently: apartment location (orientation, floor, shading) and building envelope (infiltration rate, insulation).

Modeling Thermostat Usage

To isolate the effect of thermostat behavior from other sources of variability, we performed computer simulations using the observed temperature histories as setpoint schedules in whole building simulations. The modeling approach of Urban et al. (2012) was followed for this experiment. This involved first setting up a template input file that represented a typical brick building. This input file was then replicated 60 times, each referencing a unique apartment's schedule of hourly temperature observations from the field study. Then, in Energy Plus, the input files were run in a batch process to calculate energy consumption during the two-month experiment period. The results took approximately between 15-20 minutes to run on an ordinary laptop computer. We used an actual meteorological year weather file for performing this simulation (Weather Analytics 2012). Heating energy consumption for each schedule was then extracted for further analysis. The result was a series of unique heating energy consumption values for the modeled building, each associated with a given thermostat pattern. In addition we ran several default cases, representing the ASHRAE 90.2 program, the fixed schedule default program used in the study (Figure 1), and a constant setpoint for comparison.

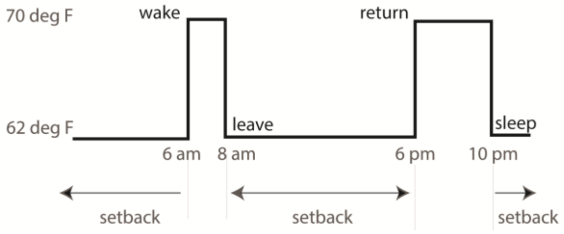

Figure 1 Default thermostat setpoint program.

Behavioral variation is unaccounted for when using fixed thermostat schedules in building simulation. This is problematic because these predictions affect the sizing of heating and cooling systems, and consequently, energy consumption. Moreover, capturing the full spectrum of behavior in models is particularly important when simulating energy saving attributable to an energy-saving thermostat.

FIELD STUDY

Methodology

The thermostat study took place in a low income housing masonry apartment building in Revere, MA during the winter of 2012 (Jan-Mar). Participants were chosen through an opt-out process, and new programmable thermostats were installed in all participating apartments. Because the building was a mixture of owners and renters, only the rental units were eligible to participate. Each unit was heated with its own gas furnace. All kitchen appliances in the units were electric, and none of the units had laundry appliances. Occupants in this building were responsible for paying their own metered gas heating and electric utility bills. Ideally, this should have provided motivation for participants to reduce or manage their heating energy consumption in this winter study.

The field study began with the replacement of the thermostats in all 82 of the participating apartments. Prior to the experiment, all units had non-programmable thermostats. Non-intrusive sensors were installed in each apartment to collect temperature and relative humidity data in 10 minute intervals and to record the times when furnace on/off state changes occurred. Upon installation, each thermostat was programmed with the correct time and with the same energy-saving default program (Figure 1). Afterwards, the occupants were free to control the thermostat in any way they desired.

The default setpoints of the thermostats had four time periods and associated temperatures for a 24-hour period during the heating season: wake, leave, return, and sleep. This default schedule is intended for a typical household that is away during the daytime hours and asleep at nighttime. Note that these default values are similar to the ASHRAE 90.2 prescriptive requirements for modeling thermostat energy use. If users had not touched their thermostats and instead had relied on the default schedules, then the standard modeling assumptions would have been valid.

At the termination of the experiment, we entered the apartment units and read the schedules directly from the thermostats to determine if the occupants had manually altered the default programs.

Experimental Results

Observed heating energy consumption varied 10-fold throughout the apartments in this experiment, ranging from less than 20 therms to over 200 therms. Variation, as shown in Figure 2, depends on orientation and floor level. Upper floors appeared to consume less energy than the lower floors, and apartment orientation had less of an impact on energy consumption.

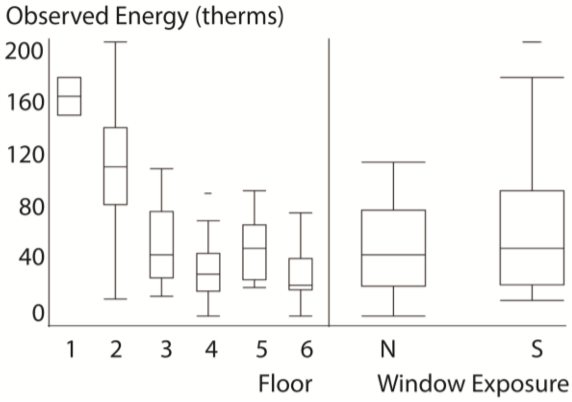

Figure 2 Observed heating energy consumption.

Moreover, less than five apartments appeared to have successfully reprogrammed their thermostat. Instead, occupants used thermostats primarily with a long-term hold, or with a combination of schedule and manual overrides. It appears that some units controlled their temperature by opening windows. Although window operating behavior was not the focus of this study, it does play a strong role in heating energy consumption.

We organized the apartments into four groups based on how they appeared to use their thermostats. First we classified apartments based on whether they used *fixed setpoints* (permanent hold pattern) or a *schedule* (program or predictable manual routine) as the basis for their control scheme. Next, we classified apartments based on the frequency of manual overrides (temporary overrides, manual adjustments, turning the furnace off, or opening windows), labeling them as *frequent* or *infrequent*. We visually inspected the graphs in Figure 4 to classify the units into these four behavioral categories.

Table 1 shows the average temperature and gas consumption throughout the experimental period. Temperatures for all groups averaged similarly around 72°F; however the standard deviations differed substantially among the groups. Notably, the units with frequent manual adjustment had higher deviation in average unit temperature. Observed energy consumption was similar among all groups but one. Those using setback schedules with infrequent overrides consumed 65% less energy on average as the other groups. Only 25% of the units operated their schedules in a way that was most consistent with energy savings.

Table 1 Observed room temperatures and gas usage

	Temp °F		**Therms**	
Pattern / Override	**Avg.**	**St. Dev.**	**Avg.**	**St. Dev.**
Fixed / Infrequent	72.2	2.8	77	58
Fixed / Frequent	73.0	4.7	76	32
Schedule / Infrequent	71.0	2.6	26	14
Schedule / Frequent	71.3	5.0	75	56
All Units	**71.9**	**4.0**	**64**	**48**

Figures 3-4 show the observations of each apartment unit over the duration of the experiment from January to March 2012. On the vertical axis is the recorded indoor temperature, and on the horizontal axis, time of day is represented from midnight to midnight. The horizontal bands represent the default thermostat program values of 70°F and 62°F. Total heating energy consumption is shown for each apartment.

Each graph in Figures 3 shows the entire temperature history of one apartment unit, with an overlay of temperature history for a sample day. For the fixed setpoint apartments (left side), we observe temperatures in a relatively narrow band. For the scheduled setpoint apartments (right side) we observe more diurnal variation. In both cases, as you move down the chart, the variability increases due to increasing frequency of manual overrides and erratic behavior. Note that most units operated well above the default thermostat program temperature range.

Darkened boxes below the graphs indicate when the heating system was on. This gives a feel for how quickly the temperatures respond in a given unit and shows when temperatures may rise due to other factors, such as solar gains. In many units, the furnace is on nearly all day long, while others cycle frequently or intermittently.

Figure 4 shows the same data with a black overlay of all temperatures that were recorded while the HVAC system was on. This overlay provides clear insight as to how a particular unit was controlling its HVAC system. Bands where black dots are absent indicate times when a schedule is in place. Dark bands indicate temperatures and times that heating was required. Isolated black dots indicate times when manual overrides or window operation took place.

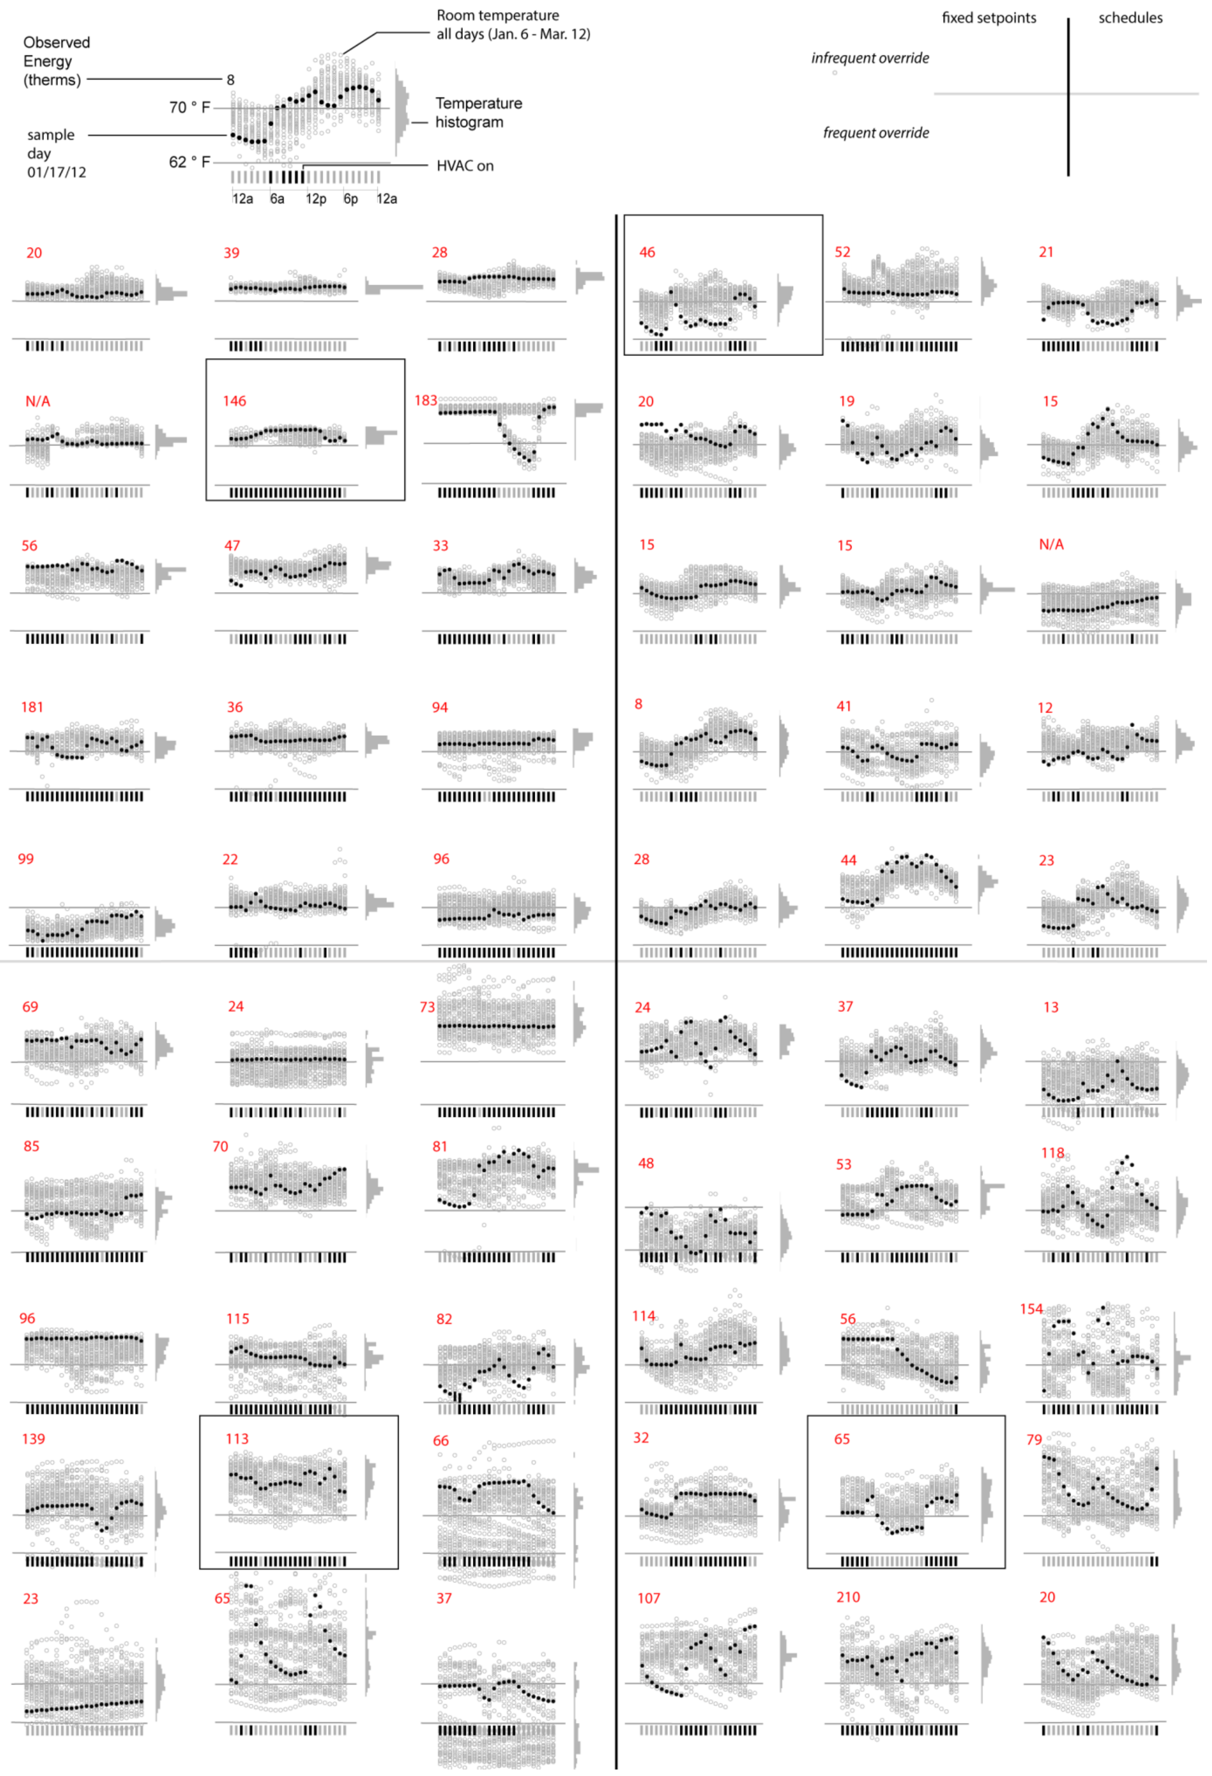

Figure 3 Temperature histories for all apartments with a sample day overlaid in black. Heating activity during the sample day is indicated by black boxes below the temperature plot.

Figure 4 Temperature histories for all apartments.
Black dots indicate the times that the heating system was on.

The wide variety in behaviors among units and the resultant heating response is abundantly clear. Most units kept their temperatures well above the default, energy-saving thermostat settings. Four selected apartments (boxed graphs in Figures 3 & 4) represent each of the four behavioral regimes.

First, the upper-left graph (146 therms) shows preference for fixed setpoints with infrequent overrides. Occupants maintained a narrow, warm temperature band throughout the winter without evidence of setbacks, and the heating system was on for nearly all of the sample day. In Figure 4 the black dots cover the spread of gray dots, indicating that heating was engaged throughout the observed temperature range. Gas consumption was high.

Next, the lower-left graph (113 therms) shows a home that appears to prefer a variety of temperatures, and seem to manually adjust the temperature up and down using a long-term hold. Again, the black dots cover the gray dots, indicating no regular pattern of setbacks. The graph just below it (65 therms) shows many individual, disconnected black dots outside of a narrow band of high temperatures. These isolated dots suggest rapid changes in indoor temperature and provide evidence for window operation, short-term manual overrides, or other temporary heat sources.

In contrast, the upper-right graph (46 therms) shows an apartment that is clearly using a setback schedule. On the sample day, the HVAC system is off at night and during the daytime. The same pattern is seen in Figure 4, with heating occurring rarely during the night or mid-day. The graph to its immediate right (52 therms) also clearly indicates daytime setbacks (but not nighttime setbacks). That only about 25% of participants routinely used setbacks is consistent with prior studies.

Finally, the lower-right graph (65 therms) shows evidence of a daytime setback schedule, but with some random deviation. Night temperatures are irregular, suggesting frequent manual adjustment.

Units remained warm even when heating systems were unused. Units rarely heated when indoor temperatures fell below 68°F, likely because of heat transfer between units and masonry's thermal mass. Units in the bottom two rows (37 and 66 therms) spent much time below 62°F, likely due to window opening during an extended period of vacancy.

SIMULATION

Methodology

The purpose of this exercise is to help understand the uncertainty that variable temperature setpoints may impose on building energy consumption. An energy model for a typical midrise apartment building (DOE 2012) was modified and used to assess this effect. We simulated the building 60 times, each time with a different variable setpoint schedule based on observed data, to generate a unique heating energy result for each input schedule. Unlike the field study, where each apartment operated its own thermostat independently, these simulations control all apartments in the building according to one schedule. This was done for simplicity and only to show the relative magnitude of behavioral effects on whole building energy consumption. In reality, there is some heat transfer between rooms, which would add another layer of complexity to the analysis.

We modified the DOE models with typical brick wall envelope characteristics, determined from BEopt (NREL 2012), to more closely represent the building in the study. The building wall and window material specifications are noted in Table 2. Simulations were run with infiltration at 1 air change per hour (ach), per the defaults of the reference model.

Table 2 Simulated wall and window properties

Wall Properties	Brick	Plywood	Cavity	Gypsum
Thickness (m)	0.10	0.01	0.09	0.01
Conductivity (W/m-K)	0.79	0.12	0.05	0.16
R-Value (SI)	0.13	0.11	1.65	0.08
R-Value (IP)	0.73	0.63	9.35	0.45
Window Properties				
U-Value IP/(SI)	0.48	(2.73)		
SHGC	0.40			

An actual meteorological year weather file (Weather Analytics 2012) was used to ensure that behavior caused by weather would be correctly addressed in the simulation. For example, if there was a warm day with a lot of sunshine, this may result in occupants adjusting their thermostats or opening their windows.

We ran a series of energy analyses using Energy Plus to review influence of behavioral use of changing patterns on energy use. Three schedules were run for each of the 60 data sets (3x60=180 simulations), along with 3 separate runs for the default thermostat program, ASHRAE program, and a constant temperature of 68°F.

1. **Actual schedule**: For each apartment, we generated an hourly room temperature schedule by averaging the 10 minute data collected by the temperature sensor. These hourly schedules were applied directly as set point schedules into Energy Plus.
2. **Maximum value**: For each apartment, we found the maximum hourly temperature throughout the experiment and used this as a fixed setpoint.
3. **Minimum value**: For each apartment, we found the minimum hourly temperature throughout the experiment and used this as a fixed setpoint.

By running these three cases, we can appreciate what level of detail is required when modeling realistic variability. More specifically, by running bounding cases (maximum vs. minimum constant setpoint temperatures) we can get an appreciation of the expected magnitude of variation. Furthermore, we can see how simulations results, based on observed temperature histories, may differ from those based on standard modeling assumptions.

A script was used to apply each of the schedules input templates run using batch routines within Energy Plus. After the simulations were complete, another script extracted the energy use for the entire building and the air temperature in an apartment. The average temperatures associated with each schedule were then calculated for analysis.

Simulation Results

Simulated whole building energy consumption based on observed schedules is shown in Figure 5. On the horizontal axis is the average temperature associated with the particular schedule. Figure 6 shows that heating energy is approximately normally distributed across the simulated schedules.

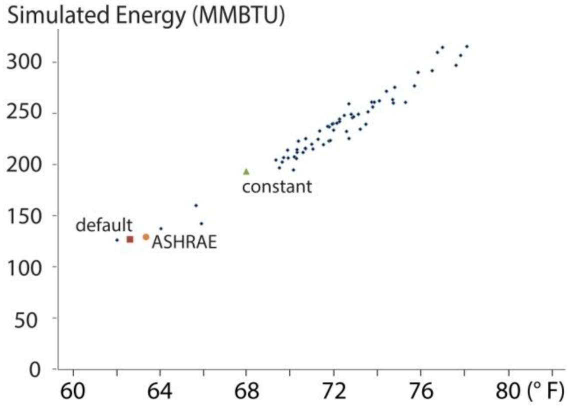

Figure 5 Simulated whole building heating energy vs. average setpoint temperature.

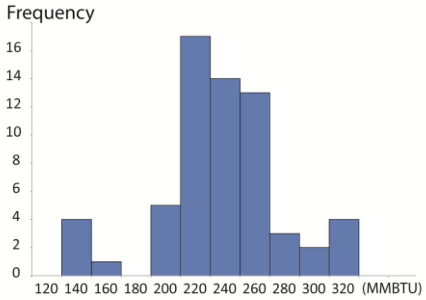

Figure 6 Simulated heating energy use histogram.

Note that the entire range of energy consumption varied by a factor of three simply because of differences in thermostat behavior. On the low end of the spectrum, we see the results from plotting the ASHRAE and default thermostat schedules. Not surprisingly, because these schedules are designed to save energy, they result in the lowest prediction of energy use. Unfortunately, they are quite distant from the points observed during this experiment. This suggests that when humans push the buttons, energy consumption goes up.

Even for a given average temperature, variation can be substantial. The width of the band is on the order of 50 MMBTU, or about 25% of the average heating consumption. This variation is expected, and can be attributed to differences in when heat is desired. An apartment with warm daytime and cold nighttime setpoints should consume less energy than one with cold daytime and warm nighttime setpoints, owing to daytime solar gains and nightly heat losses.

Among the large cluster of values to the top right of Figure 5 (above 70°F and 200 MMBTU), there is still significant variation. Predicted energy consumption varies from about 200-300 MMBTU, a factor of 1.5 across the majority of schedules. This has major implications for correctly sizing equipment in homes and for predicting the energy savings attributable to programmable thermostats. If we continue to assume that people are using idealized thermostat programs, we will continue to make poor conclusions about future energy consumption and savings.

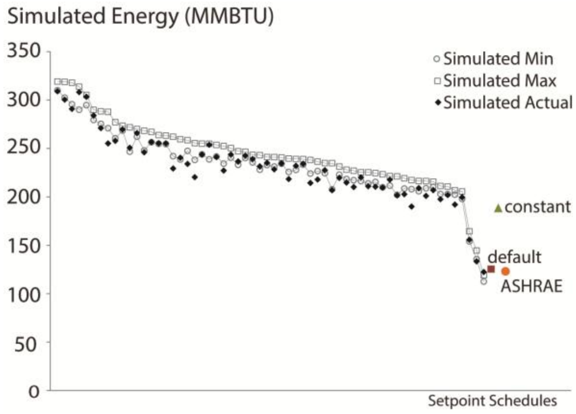

Figure 7 Simulated whole building heating energy based on observed temperatures.

For each observed schedule, Figure 7 shows the results of using the three simulation methods (constant maximum temperature, constant minimum temperature, and actual observed variable schedule). In many cases simulating the observed variable schedule yielded energy consumption estimates somewhere between the bounding cases. Sometimes, however, the actual setpoint history resulted in less energy consumption than if the thermostat had been kept at the minimum setting throughout the duration of the experiment. This is counterintuitive, but may be related to the thermal mass effect of the brick building and its ability to retain and release heat. In such cases a dynamic schedule may result in less energy consumption than a building with a constant but lower setpoint. For instance, if substantial solar gains occur during the daytime vacancy period, the building may store and release heat during the night, thereby reducing overall consumption. Further investigation is required to test the accuracy of this approach using calibrated building simulations.

Validation

Since this modeling approach is based on using existing features available in previously-validated simulation tools, the authors expect results to behave, however, modeling units with the rapid and large setpoint variability could generate unpredictable results. To validate this behavioral model against the experimental data would require more inputs than were available in this study. Developing a geometrically and physically detailed and calibrated model was outside the scope of this project, and doing so would require detailed knowledge of historic consumption, envelope characteristics, HVAC system specifications, and historic setpoint schedules for all units. Future research is required to investigate the accuracy of this modeling approach and its limitations with existing simulation tools.

We did investigate the sensitivity of results on several key parameters. Unsurprisingly, assumptions about air tightness and insulation levels had a strong impact on energy consumption results: when shifting from 1.0 ach to 0.1 ach, the simulated energy consumption curve in Figure 7 retained its shape but shifted downwards. We also repeated the simulations using typical weather files in place of actual weather data. The simulations that used typical year weather data for Boston, MA and Albuquerque, NM both produced similar patterns of energy use as did the actual weather file, again with results shifting in magnitude. This suggests that the energy implications of varying schedules may not depend strongly on specific day-to-day weather conditions.

CONCLUSIONS

Through field evidence and simulation, we have demonstrated the importance of modeling behavioral variation in thermostat settings when conducting building energy simulations.

First, the standard modeling assumptions for default thermostat setpoints appear to be overly aggressive and require a second look. Occupants in this study preferred temperatures that were far warmer than those prescribed by the ASHRAE 90.2 standard.

Second, the variation in observed heating energy consumption was significant among apartments. Energy usage varied more than ten-fold in field observations, suggesting a strong behavioral component, though factors such as building orientation and floor level were contributing factors.

Third, when the observed temperature schedules are applied to a whole building simulation, the variation is still large and significant. Simulated heating energy consumption varied by a factor of three when applying the observed schedules to the model.

Finally, we have presented four behavioral categories that represent the wide variety in temperature control: fixed setpoints vs. schedules and frequent vs. infrequent manual overrides. The group using thermostats with setback schedules and infrequent manual overrides consumed 65% less energy than all other categories, but represented only about 25% of the sample.

Modeling behavioral variability in thermostat schedules can be done readily using existing building energy modeling tools. More field studies are required to develop better understanding of the range of occupant behaviors and their underlying causes.

ACKNOWLEDGEMENT

The authors would like to thank the Building America program for its support for the experimental portion of this project.

REFERENCES

ASHRAE. 2007. ANSI/ASHRAE Standard 90.2-2007. Energy Efficient Design of Low-Rise Residential Buildings.

DOE. 2012. Commercial Reference Buildings. U.S. Department of Energy. http://www1.eere.energy.gov/buildings/commercial/ref_buildings.html

Kempton, W. 1986. Two theories of home heat control. Cognitive Science. 10(1986) 75-90.

Malnick, T., N. Wilairat, J. Holmes, and L. Perry. 2012. Destined to Disappoint: Programmable Thermostat Savings Only as Good as the Assumptions about Their Operating Characteristics. ACEEE Summer Study on Energy Efficiency in Buildings.

Meier, A., C. Aragon, B. Hurwitz, D. Mujumdar, D. Perry, T. Peffer, and M. Pritoni. 2010. How people actually use thermostats. Center for the Built Environment, U.C. Berkeley. Jan.

Meier, Alan, Cecilia Aragon, Therese Peffer, Daniel Perry, Marco Pritoni. 2011. Usability of Residential Thermostats: Preliminary Investigations. Building and Environment. 46(2011)1891-8.

Nevius, M.J. and S. Pigg. 2000. Programmable thermostats that go berserk? Taking a social perspective on space heating in Wisconsin. Energy Center of Wisconsin.

NREL. 2012. BEopt v1.2. National Renewable Energy Laboratory. http://beopt.nrel.gov.

Peffer, T., M. Pritoni, A. Meier, C. Aragon, and D. Perry. 2011. How people use thermostats in homes: A review. Building and Environment. 46(2011)2529-41.

Sachs, Olga, Verena Tiefenbeck, Caroline Duvier, Angela Qin, Kate Cheney, Craig Akers, and Kurt Roth. 2012. Field Evaluation of Programmable Thermostats. Final Report to the Building America Research Program. U.S. Department of Energy. Dec.

Urban, Bryan, Diana Elliott, and Olga Sachs. 2012. Towards Better Modeling of Residential Thermostats. SimBuild 2012. Aug. http://www.ibpsa.us/simbuild2012/Papers/SB12_TS01c_2_Urban.pdf.

Weather Analytics. 2012. Actual Meterological Year EPW file for Revere, MA. Jun. http://www.weatheranalytics.com/.