Are Thermostats the New Energy Audits?

PREPARED BY

Vermont Energy Investment Corporation
Ethan Goldman, Nick Lange, Jake Jurmain, Greg Fanslow, Rebecca Foster

January 29, 2014

Efficiency Vermont
www.efficiencyvermont.com

ABSTRACT

Vermont Energy Investment Corporation conducted a field research study to test the hypothesis that meaningful differences in household thermal performance could be discerned by data available from a programmable communicating thermostat paired with supplemental location, weather, and other broadly available information sources. Researchers were interested in exploring whether thermostat data could be used to identify the “heat-loss rate” of buildings to identify those that would benefit most from retrofit services and in the potential to isolate behavioral energy consumption from the energy used by the building itself, in other words to separate operational and asset energy use.

Data were collected from 13 thermostats in homes and small businesses in the state of Vermont. A thermal performance rating was calculated for each thermostat using a novel algorithm, which was strongly correlated with a linear regression. Results from this small sample show meaningful differences in whole-house thermal performance that are discernible based on thermostat data and a consistent relationship between ratings derived from multiple thermostats within one house. This suggests that using remote diagnostics incorporating data from thermostats may soon be a cost-effective mechanism for efficiency programs to employ. Doing so would provide useful site-specific insights and enable targeted customer engagement, messaging, and savings estimates, all of which would represent a meaningful improvement over existing generic efficiency recommendations.

Contents

Introduction 4
Approach and Methodology 5
    Research Premise 5
    Methodology 7
    Challenges Encountered 11
Results 11
    Discussion of Results 13
The Promise of Smart Thermostats 14
    What Smart Thermostats Could Deliver 15
Next Steps in the Research 15
Conclusion 16

Introduction

Vermont Energy Investment Corporation (VEIC) is a sustainable energy company that develops programs, strategies, and market interventions based on robust research and in-depth market knowledge. Efficiency Vermont—one of the comprehensive programs operated by VEIC—is charged with delivering energy efficiency services to homes and businesses throughout the state. The delivery of comprehensive residential retrofit services is complicated by the fact that 86% of the homes in Vermont are heated by unregulated fuels, including oil, propane, kerosene, and wood. As a result, accurate information on heating fuel usage for these homes is extremely difficult to obtain. Fuel dealers do not readily share information on customer fuel bills, and customer-provided information is often inaccurate because homeowners may heat their homes with multiple fuels, change vendors frequently, or carry over fuel from previous seasons. Even when fuel bills are available, they only provide monthly data and include other non-heating loads (e.g., hot water).

Lack of heating fuel usage information creates two key barriers to residential retrofit program delivery. First, it prevents program administrators from targeting programs such as Home Performance with ENERGY STAR® (HPwES) to high-use homes that are likely to have major improvement opportunities. Second, it presents a barrier to verifying actual energy savings after an improvement project is completed.

"Next generation," "smart," or "programmable communicating" thermostats (PCTs) have the potential to overcome both of these barriers. A PCT can wirelessly upload time-stamped set-point and ambient temperatures. These thermostats can identify optimal set-points based on actual homeowner behavior and a home's thermal performance in meeting the need for heat. As such, they address the pitfalls of traditional programmable thermostats and their suspect savings claims. Some PCTs, such as the Nest Learning Thermostat, are even able to detect when the home is occupied and can respond accordingly. This class of thermostats is expected to contribute to the $50 billion in annual energy savings from "intelligent energy," as defined by the American Council for an Energy Efficient Economy in a recent report.[1]

Interest in the capabilities of PCTs is not new; early efforts were focused on using them as a demand response tool. A 1988 paper (K. Subbarao et al)[2] described a project that assessed the thermal quality of a residential building based on short-term tests and a small number of data channels. The project utilized a multiple-regression model and a uniform method for building simulations to "to provide a realistically complex thermal model of a building and determine from short-term tests the statics as well as the dynamics of a building." The authors explicitly demonstrated how long-term performance and peak loads

---

[1] E. Rogers, R.N. Elliott, et al. "Intelligent Efficiency: Opportunities, Barriers, and Solutions." American Council for an Energy Efficient Economy. http://aceee.org/research-report/e13j

[2] K. Subbarao, JD Burch, et al. "Short Term Energy Monitoring (STEM): Application of the PSTAR method to a residence in Fredericksburg, Virginia" Solar Energy Research Institute, for US DOE, Sept 1988 http://www.nrel.gov/docs/legosti/old/3356.pdf

could be extrapolated and pointed to other important application areas, such as heating, ventilation, air conditioning, and controls.

While it yielded very interesting results, this early research required half a day's labor to set up the data collection channels and perform the audit, as well as multiple days of tests. The cost and complexity of data collection and computations sharply limited the practicality of broadly applying the findings. However, with the recent arrival of smart meter infrastructure across the country and the increasing adoption of PCTs, the cost of data collection and home diagnostic testing analysis are rapidly falling. Our research looked at how data from PCTs may one day soon be used to cost-effectively determine and communicate reliable thermal performance diagnostics in the residential market at scale. Though there are still significant challenges to address, the outlook appears promising.

Approach and Methodology

RESEARCH PREMISE

Efficiency industry professionals know that well-insulated houses cool down slower in the winter than poorly insulated houses. As reflected in Figures 1 and 2 below, the rate of cooling increases as it gets colder outside. Therefore, data from communicating thermostats might characterize the response to outside temperature for each individual house.

The purpose of this study was to investigate how PCTs can be used to understand whole-house thermal performance, with a focus on estimating whole-building heat flux (UA). UA can be measured by observing the amount of energy needed to maintain set-point at different outside temperatures, or by trending the cooling rate between cycles and when the heating system is set back at night.

This sort of measurement has been done under carefully-controlled conditions before, but it is an expensive and invasive process that involves multiple independent measurements of electric power, indoor temperatures, outdoor temperature and wind speed, and solar irradiance. The research team sought to investigate if a related but vastly easier method, namely only using data from a connected thermostat, might be accurate enough to use to estimate the UA of a home for the purpose of targeting retrofit program efforts.

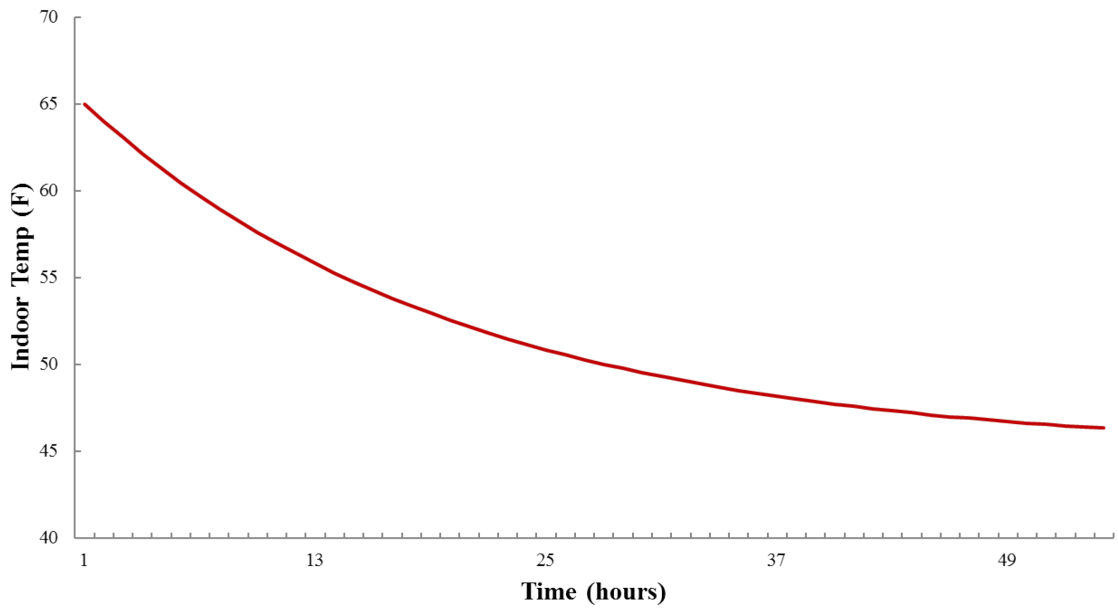

Figure 1: A simplified model of indoor heat loss based on (specific heat of air * volume * density) + internal loads – (air infiltration + ventilation)

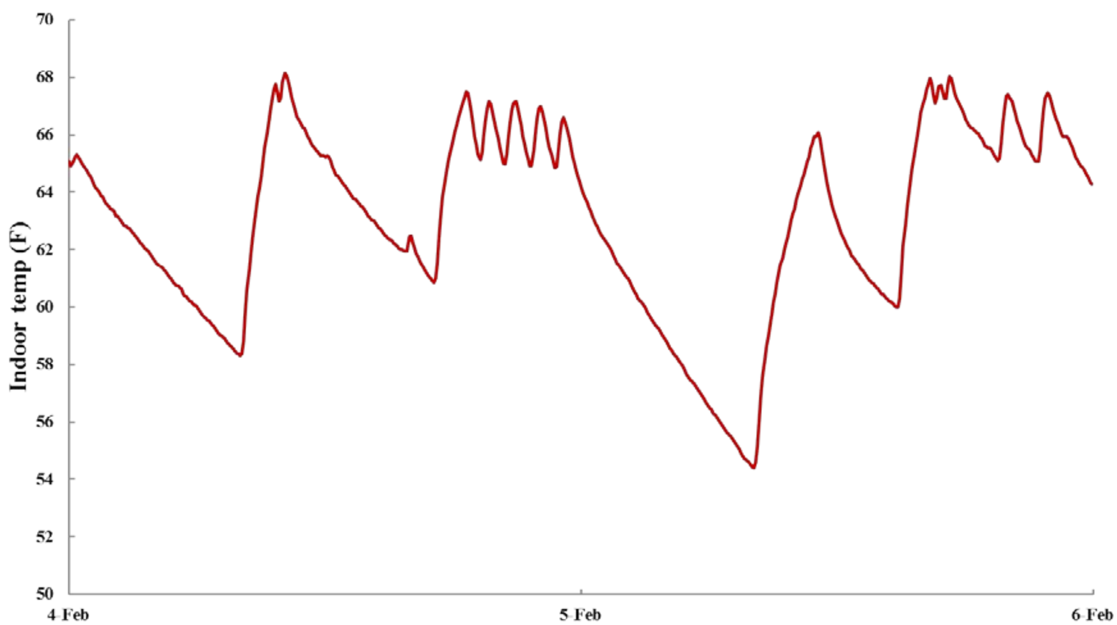

Figure 2: Indoor temperature curves show similarities to simplified case

This research sought to validate this basic premise and explore the potential to enable the energy consumption of behaviors to be isolated from the energy used by the home itself, in other words a separation of operational energy use from asset energy use. If successful, the next step in the research would be a follow-on effort to collect and analyze data from a pool of homes sufficiently large for statistically-valid conclusions about whole-house thermal performance—and resulting expected energy consumption estimates—to bring value to homeowners and efficiency programs.

METHODOLOGY

To test the hypothesis—that meaningful differences in household thermal performance can be discerned through analysis of data from a thermostat and supplemental basic location, weather, and other broadly available information sources—we developed algorithms using thermal and HVAC duty-cycle data obtained using off-the-shelf environmental data loggers as a proxy for thermostat temperature data in several homes.

The steps involved in this work were to:

1. Summarize current methods of estimating UA in existing homes.
2. Identify opportunities to leverage other related studies, including an Efficiency Vermont low-load home monitoring project, a Department of Public Service Home Performance with ENERGY STAR impact evaluation, and/or a Weatherization Innovation Pilot Program consumer behavior study of low-income households.
3. Develop data acquisition and analysis plan, including initial algorithms for estimating UA from existing datasets.
4. Acquire and install data loggers in 4-6 target households and collect data.
5. Create custom data collection software to log data remotely from participant thermostats, and log external weather data for each site from third party services.
6. Create custom data analysis software to compare the accuracy of the PCT data to the information gathered by data loggers.
7. Analyze data and make refinements to the algorithms.
8. Write a final report with conclusions.

The study database currently contains data for 13 thermostats, including 3 in a highly-efficient small commercial building. The research team collected information from homes with various kinds of heating systems, and also some multi-zone homes. Self-reported qualitative evaluations of sites includes at least one "Very Good" building shell, two "Above Average" homes, and a handful of "Average" homes. (We are evaluating low-cost ways to assess these indicators to provide a check against the results of data collection alone. This may include an online assessment tool, like Energy Savvy's 4-page survey that yields an industry-acceptable coarse evaluation metric score.)

Analytical code is able to access live data from the database, so analysis and reporting are always done based on the latest information. Figure 3 below shows a screen shot of indoor temperatures over time. The automatically identified “idealized” periods of cooling for analysis are indicated by the black brackets just above the x-axis.

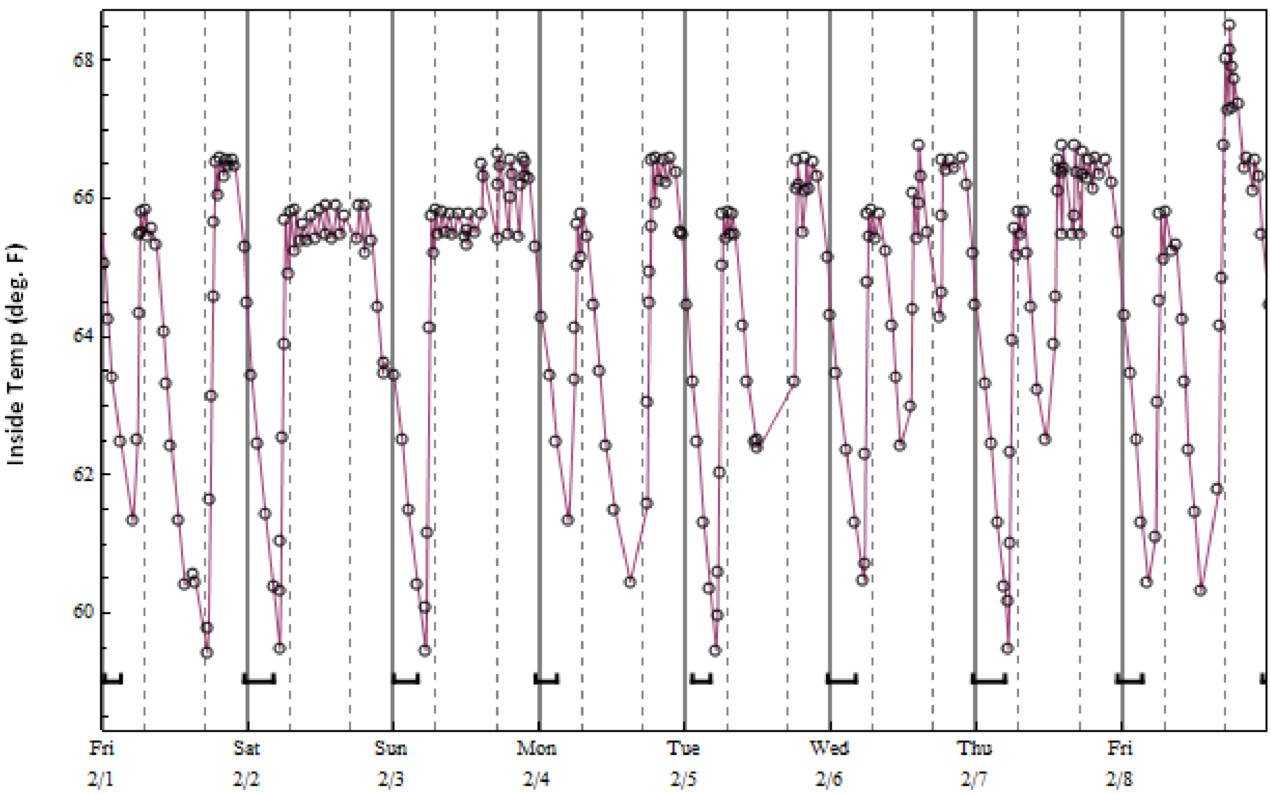

Figure 3: A screenshot showing indoor temperature vs. time, with algorithmically selected cooling periods marked.

Any thermostat's data can be selected from a list and explored. In addition, as shown in Figure 4, individual cooling periods can be selected for further analysis and detailed investigation of outliers.

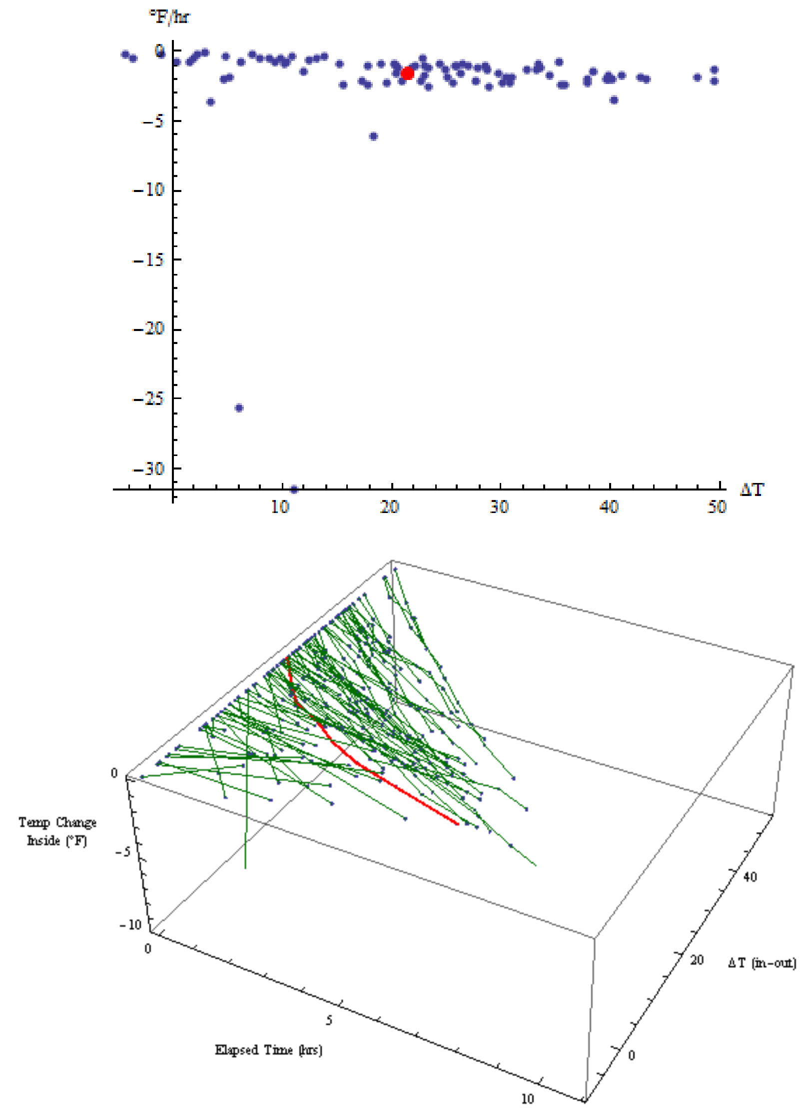

Figure 4: An individual cooling period is selected (in red) for more in-depth data exploration

Figure 5 below presents more screenshots of the analytical views for the period selected. Above is a plot of calculated heat loss rate relative to the difference in the indoor and outdoor temperatures, with linear best fit. Below is an overlay between thermostat indoor temperature data and independent temperature data loggers.

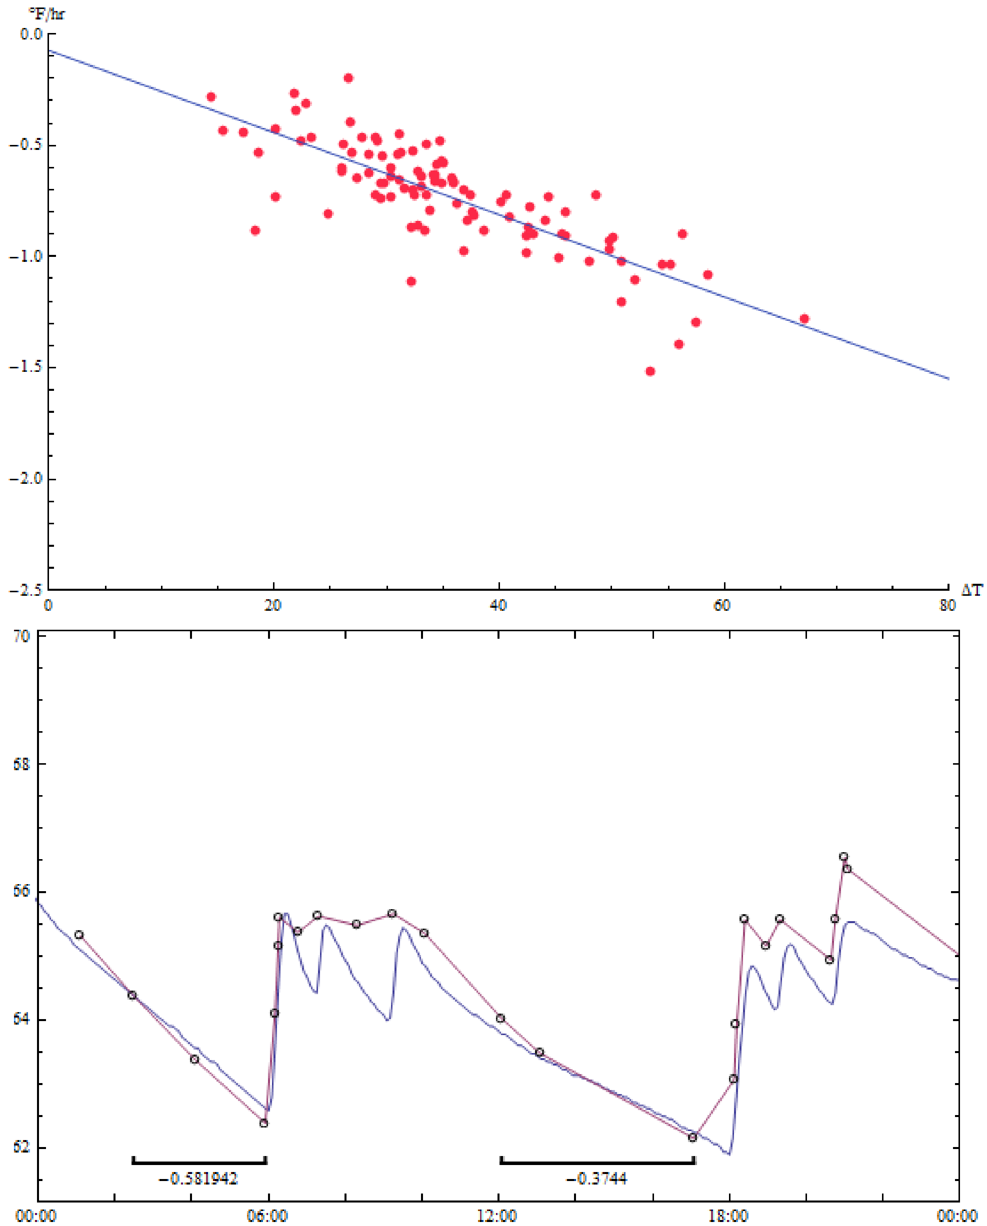

Figure 5: Additional analytical views of the data

CHALLENGES ENCOUNTERED

Data collection initially relied on small, battery-powered temperature loggers called HOBOs (from Onset Computers) that were used to gather data from a few homes. Inside and outside temperature were logged at 5-minute intervals for multiple months. Since the HOBOs had to be collected every month or so to retrieve the data, the process was not easily scalable. In order to consistently collect data from more sites over longer time periods, we began working with internet-connected thermostats as a cost-effective and easy-to-install way of remotely gathering data. While there are a growing number of options in this product category, including multiple models that offer some sort of API, we began working with the Nest Learning Thermostat because the more user-friendly installation process would likely mean more participants would be able to successfully connect their thermostats to the internet and because the API was easily accessible from anywhere, not just the home's local network (without requiring any modification of the home's firewall.)

However, while the HOBOs are dedicated data collection devices that log data on a consistent time interval (such as 5 minutes in most of our test sites), the Nest thermostat must optimize for a number of design constraints including connecting to a WiFi network using a battery that can only be intermittently charged off of the HVAC system control voltage. This seems to result in an irregular sampling frequency that is presumably intended to communicate the temperature often enough to provide a reasonably up-to-date temperature reading to remote interfaces (such as web and smart phone) using a minimum of communication energy. While the particular compression algorithm was optimized around a different use case and goals, we were still able to modify our analysis algorithms to work effectively with arbitrary-length intervals that often stretched more than an hour between samples.

Analysis in Excel was also cumbersome and time consuming. By migrating to Wolfram Mathematica, which offers not only a wide variety of mathematical functions but also a procedural and array-based programing language, we were able to develop a multi-step data analysis process that could be dynamically applied to all data collected for each location. It also allowed us to create interactive interfaces to explore the data and gain a more intuitive understanding of how the rating algorithms responded to different situations. We could then iterate on the design of the algorithms and quickly re-apply them to the entire data set, allowing many more revision cycles to improve confidence in the accuracy of the results.

Results

The data showed that homes operated in multiple modes, sometimes maintaining a set-point by cycling between brief periods of heating and cooling and sometimes drifting from a higher set-point to a lower one. Those extended cooling periods clearly showed the temperature dropping more quickly when it was cold outside, and also showed that some houses dropped more quickly than others for a particular outdoor temperature.

Excel analysis suggested that it might be possible to quantify the relationship between outdoor temperature (or more specifically, the difference between indoor and outdoor temperature, called Delta-T) and the rate of temperature decline.

We found that meaningful differences in whole-house thermal performance appear to be discernible through data diagnostics informed by a house's thermostat data. Additionally, as shown in Figure 6, we observed a consistent relationship between ratings derived from multiple thermostats within one house: the slope for each thermostat rating was similar but the x-intercept (where the rate of temperature change is zero) was consistently shifted higher for the upstairs location. Since the rough interpretation of that intercept is the balance point for the house, it makes sense that the heat flow from downstairs to upstairs would affect the balance point in a consistent fashion.

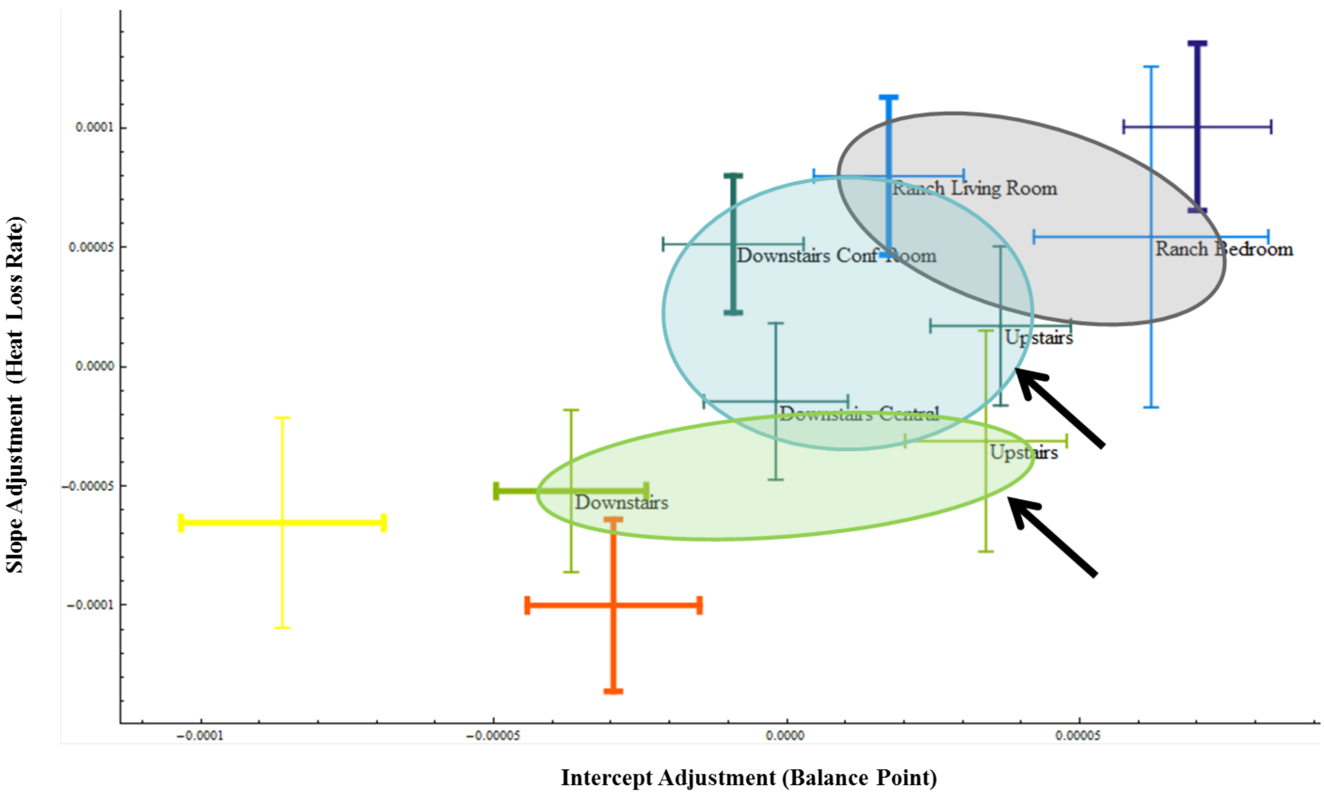

Figure 6: Distribution of calculated heat loss rates in a home

Our analytical results are also presented to illustrate the data behind the heat retention factor rating. This is in shown in Figure 7.

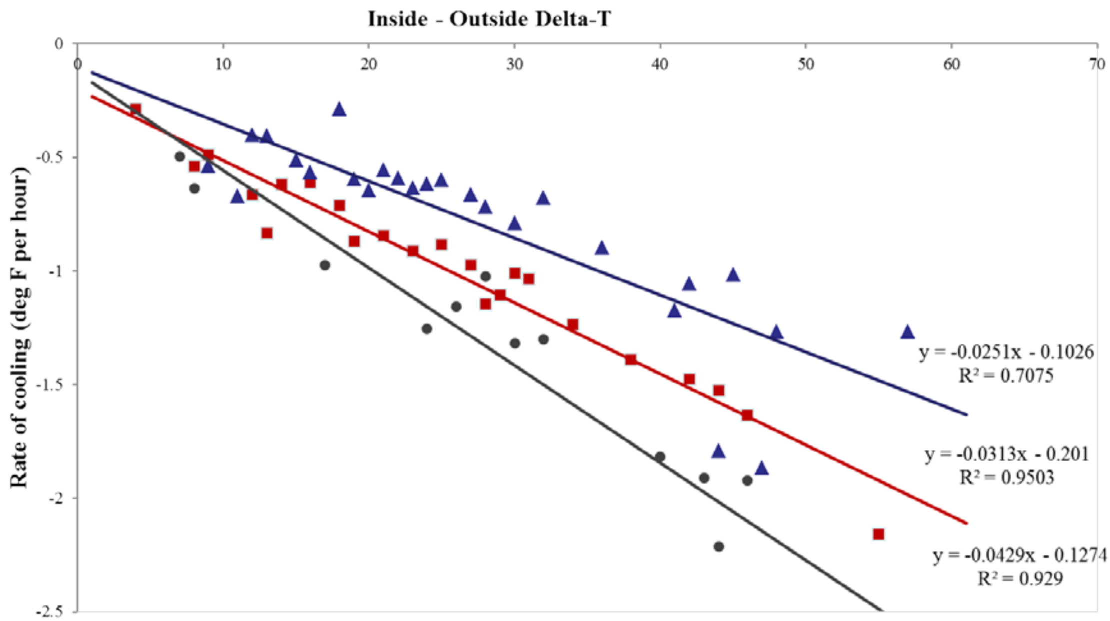

Figure 7: Degrees difference between inside and outside versus rate of cooling

DISCUSSION OF RESULTS

One challenge of this approach to measuring building heat loss is that it must isolate the effects of heat loss through the shell from all the other heat flows that are represented in the temperature trend reported by the thermostat. While the operation of the central heating system is also reported by the thermostat and can be effectively distinguished and factored out, other heating systems such as wood stoves are not controlled by any thermostat. Further, wood stoves and fireplaces have a considerable amount of mass and wood fires die out slowly, adding a non-trivial amount of noise to the temperature trend and making the isolation of building heat loss effects almost impossible. We are optimistic that, because most homes do eventually revert to central heat at least occasionally, it will be possible to detect periods that are not affected by the wood heating system and automatically limit our analysis to those periods.

Another potential confounding factor is the presence of significant thermal mass in the home, such as a masonry stove or solar thermal heat collector such as a stone floor. These features can impact the temperature trend in ways that can easily be misinterpreted as lower conductive losses. Although our rating approach necessarily combines both conductive and infiltration losses into one score, both forms of heat loss have a linear relationship with delta-T and can effectively be combined without sacrificing much accuracy (wind loading and other sources of variability in infiltration are a notable exception and a possible future topic for study). On the other hand, thermal mass will have a very different impact on the home's energy consumption and cannot be effectively combined into the same rating. For this reason, until the effects of thermal mass can effectively be dis-aggregated, those few homes with significant thermal mass cannot yet be accurately rated with this method.

Analyzing home temperature trends can also yield other insights relevant to thermal efficiency decision-making. For example, the typical estimate of 1% energy savings for each degree the thermostat setting is reduced overnight is an over-simplification. One home in our study sample had a night-time setback from 66F to 55F, and on all but the coldest nights the inside temperature never reached the setback temperature. That house has effectively reached the maximum possible savings from nighttime setbacks, as the heating system is already shutting off for the entire night and a lower set-point cannot reduce heating system run-time any further. This suggests that individualized savings recommendations could be calculated using the temperature trend data and the thermostat schedule settings, and that doing so would represent a meaningful improvement over generic recommendations. Since logically those recommendations would show the greatest savings for the first degree of setback, eventually declining to the point of zero marginal savings per additional degree of setback, it would allow users to better optimize their personal tradeoff of comfort for savings while providing more motivation to make at least a modest setback.

The Promise of Smart Thermostats

Just a few years ago, the energy efficiency industry was greeted with news that the smart grid was coming and that it would change everything. It still may, but it hasn't arrived in quite the way we envisioned: it hasn't been universally embraced and there have been some significant bumps along the way. The dawning of "big data" could be viewed in the same way. Big data enjoys the status of topic du jour and offers seemingly limitless potential. However, many consumers have yet to experience the value they have been promised, in part because the value of big data comes through its application in combination and integration with other tools and communication channels. Indeed, PCTs are a source of big data that complements the big data source from smart grid metering infrastructure. Though it may seem as if our industry is still quite far from achieving the promise of a new, smarter grid that delivers energy efficiency benefits to all customers, we may be closer than we think, and there is reason to think that PCTs will lead the way.

There are three main reasons for this. First, PCT technology is getting better and cheaper. There are more and more products available leading to a profusion of consumer choice, and increased adoption rates. Second, these products act as distributed data collection networks, meaning that efficiency programs can develop a protocol or solution to meet very specific tailored goals without waiting for an off-the-shelf solution. Third, PCTs offer programs a way to get off the ground quickly given that they are cost-effective energy-saving measures offering automatic control of heating and cooling equipment. Their data collection features may be deployed with limited added cost, especially when compared to many other stand-alone options.

WHAT SMART THERMOSTATS COULD DELIVER

Growing consumer interest in convenient and energy-saving PCTs is driving rapid distribution of communicating temperature sensors and controls into the home. Without this temperature data and control, it is hard to envision an alternative path to diagnosing the thermal performance of buildings that would not be confounded by behavioral effects in the ways that billing analyses and surveys are. The capabilities of this class of device, particularly in combination with other rapidly emerging data sources (including smart meters, smart appliances, and other technologies), may be a game-changer as a direct result of their popularity with consumers and the potency of their data.

PCTs can easily integrate with this generation of energy efficiency programs in several ways. For example, programs can benefit from cost-effective savings measures that also provide diagnostic data to drive mass personalization of messaging and savings estimates. While this is an attractive prospect, PCTs offer the most potential benefit to more innovative efficiency programs, which we believe will be needed to deliver future savings. Possibilities include using data from PCTs to:

- Enable performance-based rebates and savings claims
- Support financing programs that incorporate savings into underwriting assessments
- Improve credibility and accuracy of building labeling initiatives
- Isolate behavioral from asset-based energy consumption for more effective normative messaging and savings verification

To deliver on those possibilities, PCT-based building ratings will have to overcome several challenges. The most important are calibration with models for real-world energy consumption, including better assessment of thermal mass effects (in the heating system and in the structure), infiltration rates, and internal heat gains/losses. The overall acceptance of PCTs could be hampered by lack of common data standards and the relatively high cost and complexity of devices.

Next Steps in the Research

VEIC is excited by the potential contributions that PCTs could make to efficiency programs of the future. After gathering data from the small sample in Phase 1, we plan to continue our research and validate the promise of the initial results by calibrating the metric to energy audit building and consumption data. We plan to collect and analyze data from a pool of sites sufficiently large for statistically-valid conclusions about whole-house thermal performance and resulting expected energy consumption estimates. We have launched Phase 2 and expect to have results to share in early 2015.

The specific goals for the next phase of the research project are listed below.

1. Collect connected thermostat data from hundreds of sites with existing audit-level site information.
2. Estimate site energy performance (like annual heat load) with a mathematical model using:
    a. Thermostat-data metric developed in Phase 1 (slope regression)
    b. Basic, non-audit grade, site information (like square footage);
    c. Other elements (e.g. averages for weather and building characteristics)
3. Evaluate and compare the estimates of the model to site energy audit results.
4. Evaluate and compare estimates of this study's model to actual consumption data.
5. Assess opportunity for isolation and analysis of the behavioral component from the asset in total energy consumption through approaches derived from study results.

Conclusion

Based on our work to date, we believe that energy efficiency programs can garner significant benefits from the latent capabilities of new, intelligent technology. There is already tremendous momentum from manufacturers, consumers, and some leading utilities toward the adoption of a new generation of home (and small business) automation devices. Our experience suggests that the data from these devices, and likely also their control capabilities, can serve energy efficiency program purposes well.

Capitalizing on this promise will require several key steps. First, we must develop new data analysis techniques to understand and interpret the data these devices provide. While this is challenging, it is far more attainable than building a dedicated network of devices for energy efficiency purposes alone. Second, once programs have mass-customized performance ratings and real-time feedback at their disposal, we will be required to radically transform the way we interact with customers. Fortunately, we can look to the experience of retailers and web service providers, who have both been able to personalize and improve their services with data in order to learn how to (and how not to) communicate to customers. Third, we must tread carefully and take customer privacy concerns seriously in order to avoid backlash against the use of this new data. Effectively engaging customers and employing opt-in program designs will be important. Lastly, we must coordinate and collaborate as an energy efficiency program industry to communicate desired features and functionality to product manufacturers and standard-setting organizations such as ENERGY STAR to minimize costs and maximize benefits.

There is still a long way to go to fully realize the potential we have identified. Achieving the steps articulated above will require almost a "perfect storm" of factors to achieve a successful outcome. However, we would argue that the required efforts are justified given that the energy efficiency program industry faces ever larger savings targets and, as a whole, has struggled to identify an easily scalable way to address residential thermal efficiency. The potential for programs to play a facilitation role between customers and contractors is constrained by current mechanisms of consumer engagement and evaluation, measurement, and verification. By employing the diagnostics offered through PCTs, we see a new approach that can substantively change the strategy for interactions between efficiency programs, customers, and contractors.

In short, knowing that we may be just around the corner from significantly more relevant savings estimates, performance-based rebates, and savings guarantees should impact the way the efficiency program industry approaches investments in developing and implementing building labeling, codes and standards, and financing initiatives. More than the sum of its parts, combining thermal diagnostics through PCTs with innovative program best practices may bring exponential increases to the impacts of energy efficiency programs.