BALANCE: Budget-Driven Smart Thermostat Control

Kathryn McArdle and Christine Julien
The Center for Advanced Research in Software Engineering
University of Texas at Austin
{kate.mca, c.julien}@utexas.edu

***Abstract*—Heating and cooling (HVAC) needs account for a large percentage of residential energy usage, which has prompted development of autonomous thermostat control to reduce unnecessary HVAC usage, primarily when occupants are away. We instead consider the balance of the financial burden of HVAC costs with occupants' comfort desires while they are at home. We present BALANCE, an autonomous thermostat management system that balances a range of user comfort levels with a user-defined monthly budget for HVAC costs. We provide a simple agent that accomplishes this task in realistic simulated environments.**

I. INTRODUCTION

According to the US Energy Information Administration, HVAC systems account for 48% of US residential energy usage.[1] As such, these systems, typically controlled by a thermostat, are good targets to reduce residential energy consumption. The temperature to which the thermostat is set (known as the *setpoint*) affects how much energy the HVAC system uses, and studies have shown that many users are not setting their thermostats in energy efficient ways [1]. In this paper, we present BALANCE, an autonomous tool to help users reduce HVAC energy usage by balancing cost with comfort to regularly determine reasonable thermostat setpoints.

One approach to reducing HVAC systems' energy usage is to provide users with feedback. A review of studies on providing residents with feedback on their overall electricity usage found that feedback is most effective when it: is based on residents' actual consumption; is provided at least every day; includes a breakdown for individual appliances; and is given over a long period of time [2]. Another review of energy usage feedback studies found that studies under six months in duration typically reported greater energy savings than studies over six months, suggesting that energy savings due to receiving feedback decrease over time [3]. Both reviews found that studies involving at least daily feedback typically suffered from a small sample size, around ten households according to [2]. Thus, while both reviews report typical energy savings of 5% to 12%, it is hard to make definitive conclusions about the potential for feedback to provide reliable energy savings. Further, when considering HVAC usage, we hypothesize that daily feedback may be confusing, as the usage when one's thermostat is set to the same temperature on different days may vary, depending on the weather and occupants' activities.

Perhaps as a result of the challenges to using feedback to reduce HVAC usage, *autonomous* thermostat control has recently become commercially popular, as with Nest and Opower[2], and well-studied in multiple academic disciplines. These approaches typically reduce HVAC usage by adjusting the thermostat setting *while the house is unoccupied*, so as not to disrupt occupants' thermal comfort. In the case of heat-pump heating systems, reinforcement learning techniques have been employed to control the system's thermostat during pre-programmed times when occupants will be away, such that less energy is used during unoccupied hours and the house is at the desired comfort level when the occupants return [4]. To help alleviate the problem of peak demand on hot days, reduced-order modeling has been performed in combination with model predictive control to aggressively cool a home before peak hours, so that less energy is needed during peak hours to maintain user comfort in those hours [5]. With these solutions, the user's comfort requirements are assumed to be fixed and prioritized, and thus energy savings are limited to what can be accomplished without any user sacrifice, when the house is unoccupied. However, this approach may not always be appropriate, since, in approximately half of U.S. households, someone is at home throughout the day [1].

Therefore, to reduce residential HVAC usage, we consider alternatives to manual and feedback-based approaches, and we do not restrict ourselves to reducing HVAC usage only when the house is unoccupied. Instead, we consider a problem many households face: balancing the cost of energy bills (e.g., keeping one's energy bill below $\$100$) with the desire for certain thermal comfort (e.g., keeping one's house at $70°F$ on a hot day). Relaxing indoor temperature requirements can reduce HVAC usage and in turn lead to more desirable energy bills. Some previous work has taken this approach. In the case of real-time dynamic electricity pricing, one approach uses preference elicitation and active learning to adjust the thermostat based on the current electricity price and the user's thermal comfort desires [10]. We believe that most people would find providing user input to this system challenging, because information about how much energy usage a given thermostat setting will result in (and therefore how much it will cost the user, given the current price) is not readily available. Further, the approach assumes that the user's preferences remain consistent: given an electricity price and an outdoor temperature, the user has one desired indoor temperature. In reality, this decision may be influenced by many factors. Of primary interest to us is the consideration of other expenses the user may have. At a time when the user has just paid his annual taxes, for example, he may want to save more money (and thus want to have a lower energy bill) than when he has just received a raise at work. In this case, the user may find a range of temperatures acceptable as long as his bill at the end of the month meets his budget. To confirm this hypothesis, we conducted a survey with 40 low-income families in Austin, Texas, and we found that nearly three-quarters of participants

[1] http://www.eia.gov/consumption/residential
[2] http://nest.com and http://opower.com

are willing to relax their thermal comfort requirements by a few degrees in order to save money.

In this paper, we present an approach to autonomous thermostat control that is geared toward households in which the need to meet a certain heating and cooling cost requirement sometimes *outweighs* the occupants' desire for *ideal* thermal comfort levels; in such households, residents must balance their level of thermal comfort with the costs of heating and cooling. We hypothesize that such households are more likely to be low-income than median-income or wealthy. On average, households across the United States spend 2.7% of their income on energy bills[3], while low-income households spend 25% [11]. We believe that a lower energy bill would be a motivating factor for reducing HVAC usage in low-income households, due to the large portion of income that must be spent on energy usage. Therefore, while our approach is suitable for all income levels, we design and evaluate it specifically for low-income household settings.

We propose that households that wish to balance HVAC costs with thermal comfort need an autonomous tool to effectively achieve this balance. It is difficult for nearly anyone to mentally convert his thermostat setting to dollars on his energy bill. Factors such as the weather, occupant behavior, and the house's thermal properties influence HVAC energy usage, and the energy bill does not display usage by the heating and cooling system alone. While the user may be aware that setting his thermostat to $74°F$ on a hot day will be more expensive than $77°F$, he cannot easily quantify the extra cost. Thus, for a user that would like to set his thermostat to $74°F$ if he can afford it but is willing to set it as high as $77°F$ if not, he needs much more information than his thermostat and energy bill can provide. Even with access to the right information, it would be cumbersome and unrealistic for a user to constantly change his thermostat setting over the course of a billing cycle (typically, one month), in order to meet his budget. To that end, we present BALANCE, an autonomous system whose goal is to balance a user's monthly energy budget with a range of comfortable thermostat settings.

Our approach to automating thermostat control is inspired by a framework that manipulates Web content requested by a mobile phone [13]. This framework receives a website request from the user, estimates the user's data budget for the webpage (based on a monthly data budget and knowledge of how much has already been spent and how far into the month this request occurs), and reduces the data content of the webpage as needed. Similarly, BALANCE requests from the user a monthly budget and preferences for heating and cooling, then updates the thermostat settings throughout the month based on the preferences, usage to-date, and contextual information such as the weather. Our approach is the first to directly consider a user's monetary budget in automating thermostat control.

II. THE BALANCE APPROACH

The goal of our autonomous thermostat management system is to meet a user-provided monthly budget for heating and cooling costs, while staying within the bounds of user-provided comfort levels. Specifically, there are four comfort levels: *preferred* low and high thermostat settings when only comfort is considered; and *absolute* low and high settings when cost is also considered. For example, one might have a preferred low temperature of $67°F$, preferred high temperature of $74°F$, absolute low temperature of $62°F$, and absolute high temperature of $77°F$. More specifically, the user would like the HVAC system, at all costs, to avoid allowing the house to be colder than $62°F$ or warmer than $77°F$. The user would *prefer*, if possible given his monetary budget, to maintain the house between $67°F$ and $74°F$. We assume that all indoor temperatures within the *preferred* range satisfy the user equally; that temperatures between the *preferred* and *absolute* settings satisfy the user less as they approach the absolute temperature; and that temperatures outside the *absolute* range are unacceptable.

With the budget and temperature preference information, our general approach is to adjust the house's thermostat setting at regular time steps (for example, every hour) throughout the month, with the goal of meeting the user's budget. At the beginning of each time step, BALANCE calculates a budget for the coming time step. It also collects the current outdoor and indoor temperatures and the energy used by the HVAC system in the previous time step. With the information it has collected over time, it estimates the optimal thermostat setting for the coming time step, while staying within the bounds of the *absolute* range, and it sets the thermostat to that setting. A thermostat setting consists of a heating setpoint, which is the indoor temperature at which the heating system should turn on, and a cooling set point, which is the indoor temperature at which the cooling system should turn on.

Each implementation of this general approach will define the length of a time step, but a time step should generally be less than one day. While BALANCE makes a prediction for how much energy will be used by a given thermostat setpoint, its prediction is unlikely to be perfect, due to the significant impacts that weather and occupant behavior can have on the indoor temperature. Thus, for any time step length, there is a risk of the actual cost of the time step exceeding or being less than the time step's budget. With a time step greater than a day, this difference could become quite large, and BALANCE has to wait until a given time step is over before it can "recover" by updating the thermostat setpoint. However, with a very short time step (under two minutes, for example), the HVAC system may not have the opportunity to bring the house's indoor temperature to the desired temperature before the thermostat setpoint is updated again. In this case, the energy usage recorded would reflect much more significantly on the current indoor temperature than on the thermostat setpoint, making prediction more challenging. Therefore, the time step length should balance the computational costs of determining the appropriate thermostat setting with the expected frequency of changes in weather and occupant behavior and with the expected HVAC operation time.

A. Algorithm

Given the above generic description of our BALANCE approach, there are two calculations that must be made at every time step, the details of which we have not yet covered: the time step's budget and an estimate for the optimal thermostat setting. There are many ways these calculations could be made, and we now present our specific implementation, which we

[3]http://www.eia.gov/todayinenergy/detail.cfm?id=10891

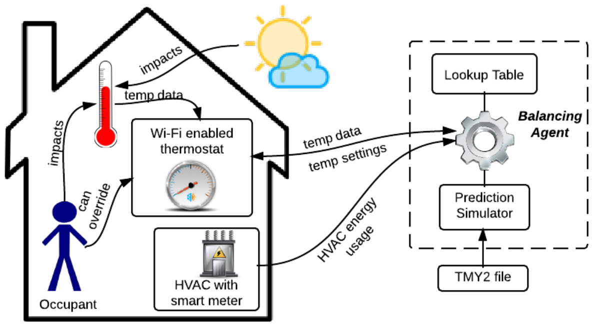

Fig. 1. Depiction of *BalancingAgent* in a real-world setting

Key Tuple			Value Tuple	
outdoor temp (°F)	*indoor temp (°F)*	*energy (kWh)*	*heating setpoint (°F)*	*cooling setpoint (°F)*
85	76	1.25	67	74
85	76	0.75	67	76
77	74.2	1.13	67	75

TABLE I. SAMPLE LOOKUP TABLE

Key Tuple			Value Tuple	
outdoor temp (°F)	*indoor temp (°F)*	*energy (kWh)*	*heating setpoint (°F)*	*cooling setpoint (°F)*
85	76	1.31	67	74
85	76	0.75	67	76
77	74	1.13	67	75

TABLE II. SAMPLE LOOKUP TABLE, UPDATED

call *BalancingAgent*. We recognize that the HVAC energy cost any agent can effect is upper-bounded by the cost of using the *preferred* temperatures for the whole month and lower-bounded by the cost of using the *absolute* temperatures for the whole month. Therefore, we use the *BalancingAgent* to evaluate the general idea of autonomous energy budgeting.

We depict the real-world setting of the *BalancingAgent* in Figure 1. The user directly provides the thermostat with his budget and temperature preferences. When the system is first set up and any time the user updates his preferences, the thermostat relays the user's preferences to the *BalancingAgent*. The weather and the occupants' behaviors impact the indoor temperature, which is detected by a thermometer on the thermostat. The weather also impacts the outdoor temperature, which is detected by the thermostat either via a connection to an outdoor thermometer or by querying a nearby weather station via the internet. At every time step, the thermostat relays the current indoor and outdoor temperatures to the *BalancingAgent*. The HVAC system is outfitted with a Wi-Fi enabled smart meter that measures its energy usage. At every time step, the smart meter relays (via a Wi-Fi connection to the thermostat) the energy used in the last time step to the *BalancingAgent*. Thus, at every time step, the *BalancingAgent* receives the current indoor and outdoor temperatures and the energy used by the HVAC system in the previous time step. It uses the energy used in the previous time step to calculate the next time step's budget and searches for the budget and temperatures in a lookup table to determine the appropriate thermostat setting. The budget calculation and the lookup table are described in detail below. In the case that the given budget and temperatures are not present in the lookup table, the *BalancingAgent* uses a simulator, also described below, to determine the thermostat setting. Once the *BalancingAgent* has determined the setting, it relays the decision back to the thermostat, which updates the setting accordingly.

The most basic and novel aspect of our autonomous thermostat agent is its calculation of the monetary budget for the coming time step. We do this hierarchically. First, at the beginning of each day the *BalancingAgent* calculates a budget (in dollars) for that day, b_{day}. The day's budget is simply the ratio of the user's monthly budget (b_{month}) that has not been used to the number of days left in the month:

$$b_{day} = \frac{b_{month} - b_{month,used}}{N_{days,month} - N_{day}} \quad (1)$$

The portion of the month's budget that has been used ($b_{month,used}$) is easily calculated by multiplying the energy used (in kilowatt-hours, kWh) at each time step by the electricity price (p, in $\frac{\$}{kWh}$), which is available from the utility company. Second, the time step's budget (b_t, in dollars) is calculated as the ratio of the day's budget that has not been used to the number of time steps left in the day. Given the time step's budget in dollars, the *BalancingAgent* calculates the time step's "energy budget" by dividing the daily budget by the electricity price.

To calculate the thermostat setting for a particular time step, the *BalancingAgent* first assesses if the energy budget is negative, as when the daily or monthly budget has been exceeded. If it is, the *BalancingAgent* returns the absolute low and high temperatures as the heating and cooling setpoints, respectively. Otherwise, the *BalancingAgent* employs a lookup table that maps the key tuple (*outdoor temp*, *indoor temp*, *energy budget*) to a value tuple (*heating setpoint*, *cooling setpoint*). This table begins empty when the system is first set up for the house, and it is populated as the *BalancingAgent* experiences different temperature-energy-setpoint combinations. A sample lookup table is shown in Table I. If the key tuple is in the lookup table, the *BalancingAgent* uses the corresponding value tuple as the setting. For example, if the current outdoor temperature is $85°F$, the current indoor temperature is $75°F$, and the time step's energy budget is 1.25 kWh, the *BalancingAgent* would select $67°F$ and $74°F$ as the heating and cooling setpoints, respectively. If the key tuple is not in the lookup table, the *BalancingAgent* runs simulations (described below) for multiple potential thermostat settings to find one that matches the budget. At the end of each time step, the *BalancingAgent* updates the table with the key tuple (*outdoor temp at start of time step*, *indoor temp at start of time step*, *actual energy used during time step*) and value tuple (*heating setpoint used*, *cooling setpoint used*). Continuing with our example, assume that at the end of the time step, the HVAC smart meter reported that the HVAC system actually used 1.31 kWh in the previous time step. The first row in the lookup table would be updated as displayed in Table II.

B. Prediction Through Simulation

To estimate an optimal thermostat setting when the current key tuple is not present in the lookup table, we use the

GridLAB-D[4] simulator. This simulator models the response of a single-family house over time to many parameters. Constant parameters include the size of the house, type of heating and cooling system, and thermal efficiency, among many others. GridLAB-D takes as input a Typical Meteorological Year (TMY2)[5] file that specifies typical weather behavior for a particular US city. Each TMY2 file contains meteorological data collected in the given city over the course of several years, and from this data the simulator approximates typical climate conditions over the course of one year.

We use the GridLAB-D simulator within our *BalancingAgent* as a predictive tool. For any time step that the *BalancingAgent* must simulate for prediction, we use the same constant parameters, which approximate the actual house in question. Using the current indoor and outdoor temperatures as input, the *BalancingAgent* runs a series of "predictive simulations" for the time step and varies the thermostat settings within the range of temperatures that fall between the user's preferred and absolute temperatures. Each simulation is run for the duration of the time step, and the *BalancingAgent* compares the energy used by the HVAC system for each simulation to the energy budget. We select the thermostat setting whose energy usage most closely matches the time step's energy budget without exceeding it. For example, assuming the lookup table in Table II, when we encounter a situation in which the current outdoor temperature is $89°F$, the current indoor temperature is $76°F$, and the time step's energy budget is 1.28 kWh, we find that this key tuple is not in the table. In this situation, the *BalancingAgent* runs a predictive simulation, specifying the given current outdoor and indoor temperatures, and using $67°F$ (the user's preferred low temperature) as the heating setpoint and $74°F$ (the user's preferred high temperature) as the cooling setpoint. After this predictive simulation completes, the *BalancingAgent* records the energy used by the HVAC system and compares it to the energy budget of 1.28 kWh. Assume the HVAC system used 1.45 kWh in this predictive simulation. This energy usage is greater than the budget, so the *BalancingAgent* must adjust the setpoints and run another predictive simulation. Because the house must be cooled, the *BalancingAgent* increases the cooling setpoint by a degree, to $75°F$, and runs another predictive simulation. This process continues until the *BalancingAgent* finds a setpoint that results in less energy use than the energy budget or reaches the user's absolute high temperature of $77°F$.

In this situation, one alternative to running predictive simulations would be simply to use the "cheapest" option that is acceptable to the user: set the heating and cooling setpoints to the user's absolute low and high temperatures. However, this approach would overemphasize the user's budget, instead of balancing it against comfort, as BALANCE aims to do.

In the predictive simulations, we specify an outdoor temperature that matches the outdoor temperature at the *start* of the time step, but, within the simulation, the change in the outdoor temperature over the course of the time step is governed by data in the TMY2 file for the city in our simulation. Therefore, the change in temperature over the course of the time step is *typical* for the particular time step being simulated but may not be an exact representation of what happens for this specific instance. This potential discrepancy motivates the use of reasonably small time steps for the predictive algorithm.

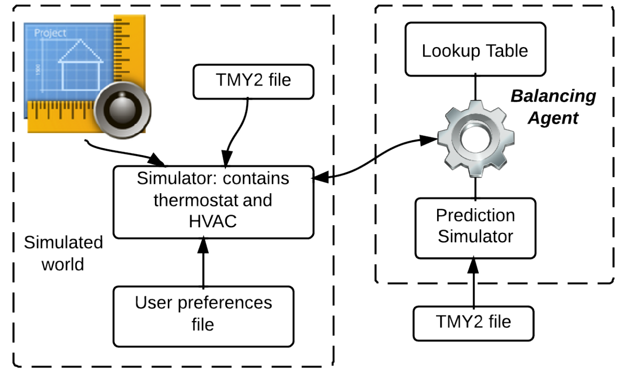

Fig. 2. Depiction of *BalancingAgent* in a simulated setting

III. EVALUATION

Because evaluating our approach in a real-world setting would take a full year to gather results on our agent's performance in all four seasons, we use a simulator instead, and we leave a real-world deployment to future work.

A. Experimental Setup

To evaluate the *BalancingAgent*, we use the GridLAB-D simulator. The constant simulation parameters we use are a floor area of 1000 square feet and an electric central heating and cooling system. Other parameters that would affect the home's thermal response are automatically calculated by the simulator, based on the two we provide. To model realistic outdoor temperatures, we provide the simulator with the TMY2 file for the desired city in our simulation.

Figure 2 depicts the setting of the *BalancingAgent* in the simulation. As compared with Figure 1, there are some differences to note: the thermostat and HVAC have been replaced by the simulator; the user has been replaced by an input file containing realistic user preferences; the house has been replaced by values of the house parameters; and the weather has been replaced by a TMY2 file. The *BalancingAgent* itself remains the same. Many of these replacements are typical of simulations, and the simulator itself has been well-developed and exercised for real-world modeling (e.g. [4], [14], [15]).

At first glance, our use of the simulator for both the overall simulation *and* the prediction mechanism poses one additional challenge: for a given time step, the agent is basing its prediction on the exact situation that will occur during that time step in the overall simulation. In a real-world deployment, the simulated time step used to make predictions would likely not exactly match the conditions in the real world: the agent would not be able to specify the weather over the course of the time step, for example, because the agent will not be able to perfectly predict the weather. However, we note that we are using the prediction only to determine the thermostat setting, which must be rounded to whole numbers for use with a thermostat. In the prediction phase, the agent seeks the (integer-valued) thermostat setting whose energy use most closely

[4]http://www.gridlabd.org
[5]http://www.nrel.gov/rredc

Name	Min. Temp. ($^\circ F$)	Pref. Low Temp. ($^\circ F$)	Pref. High Temp. ($^\circ F$)	Max. Temp. ($^\circ F$)
PrefersCool	62	67	74	77
PrefersWarm	67	72	79	82

TABLE III. TEMPERATURE PREFERENCES USED IN EVALUATIONS

meets the budget without exceeding it; this rounding introduces an important level of approximation to the prediction. The integer-valued thermostat setting will usually not correspond to energy use matching the *exact* energy budget, because of the rounding that must take place. Thus, knowing the exact energy that will be used in a time step is not a significant advantage over having an approximation for the predicted energy usage. Additionally, the agent will have access to a weather forecast, which – while not a perfect prediction – is fairly reliable for short periods of time. Therefore, we conclude that our use of the same simulation for prediction and evaluation is justified.

We run simulations for Austin, Texas, for the months of January and August. We use $\frac{\$0.08626}{kWh}$ as the electricity price in Austin during summer months (June through September), which is the average summer price based on the tiered pricing that the City of Austin utility uses[6]. Similarly, we use $\frac{\$0.04580}{kWh}$ for the Austin winter electricity price (October through May). We evaluate our *BalancingAgent* with two sets of temperature preferences, as shown in Table III. These preferences are based on results from the survey discussed earlier.

To measure the performance of our *BalancingAgent*, we developed two baselines for comparison. The *LowestCostAgent* favors cost savings by setting the heating setpoint to the temperature used for the *BalancingAgent's* minimum temperature and the cooling setpoint to its maximum temperature. The *HighestComfortAgent* favors comfort by using the *BalancingAgent's* preferred low and high temperatures for the settings. We note that the cost of the *LowestCostAgent's* energy usage represents the minimum possible cost any agent could obtain while staying within the user's absolute range of temperatures. Similarly, the cost of the *HighestComfortAgent's* energy usage represents the maximum cost any agent should cause, given the user's preferred temperatures. Therefore, we evaluate our *BalancingAgent's* ability to meet various budgets that fall between these minimum and maximum costs.

We choose a time step of one hour: every hour, we run the *BalancingAgent* algorithm and instruct the simulator to set the thermostat to the resulting setting. At the end of each time step, we record the energy used by the heating and cooling system. The current indoor and outdoor temperatures that are recorded at the beginning of each time step correspond to the respective temperatures at the end of the previous time step.

B. Results

We measure the success of our algorithm in two ways: how closely it meets the user's monthly budget and how little it deviates from the user's preferred comfort range, i.e., the temperatures between his preferred low and high temperatures. To measure the month's budget success, we multiply the month's total HVAC electricity usage by the electricity price

[6]http://austinenergy.com/wps/portal/ae/Residential/Rates

Month	Temp. Prefs.	Agent	Budget ($)	Cost ($)	% Budget Deviation	% Mins. Outside Range
Jan.	*Prefers Cool*	*LC*	–	10.86	–	76.74
		B	13.00	13.21	1.63	38.90
		B	15.00	13.88	−7.47	21.35
		HC	–	16.33	–	0.09
	Prefers Warm	*LC*	–	10.33	–	73.81
		B	13.00	12.63	−2.84	33.26
		B	15.00	13.32	−11.17	23.12
		HC	–	17.00	–	0.07
Apr.	*Prefers Cool*	*LC*	–	24.02	–	94.39
		B	24.50	25.39	3.61	41.72
		B	25.50	25.62	−0.48	33.46
		HC	–	26.53	–	0.13
	Prefers Warm	*LC*	–	19.96	–	86.95
		B	20.50	21.28	3.82	41.69
		B	21.50	21.54	0.20	31.78
		HC	–	22.41	–	0.10
Aug.	*Prefers Cool*	*LC*	–	83.30	–	100
		B	85.00	85.76	0.89	59.08
		B	87.00	86.43	−0.65	48.68
		HC	–	89.39	–	0.15
	Prefers Warm	*LC*	–	73.13	–	100
		B	75.00	75.68	0.90	57.73
		B	77.00	76.27	−0.95	48.68
		HC	–	79.23	–	0.10

TABLE IV. SIMULATION RESULTS. *LC = LowestCost, B = Balancing, HC = HighestComfort*

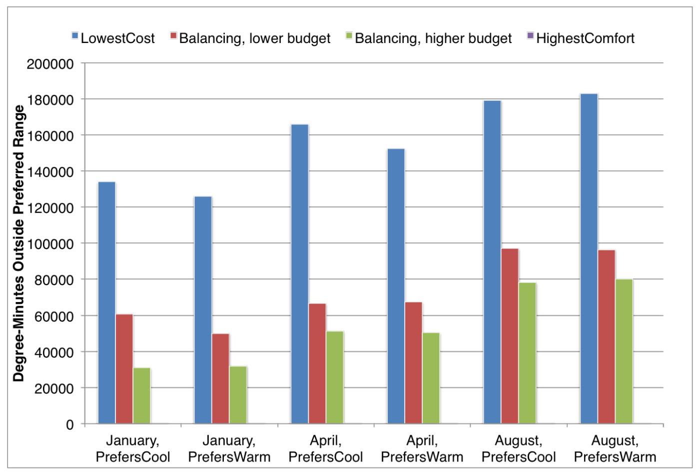

Fig. 3. Comparison of Degree-Minutes spent outside the preferred range

to calculate the total cost of the month and compare it to the provided budget. To measure the deviation from the comfort range, for each time step we record the number of minutes when the indoor temperature is outside of the preferred comfort range and the highest deviation in degrees between the actual temperature and closest edge of the range. We multiply these two factors together to estimate *degree-minutes* that are in violation of the preferred comfort range.

Table IV shows the results of our simulations, and Figure 3 displays the difference in comfort that each agent achieves in the three months. In January, the *LowestCostAgents* spent most of the month outside the preferred comfort range, while the *HighestComfortAgents* rarely exceeded it. The *BalancingAgents* spend much less time outside the preferred comfort range than the *HighestComfortAgents* while also meeting the budget. The actual cost used by each *BalancingAgent* is several dollars below the cost of the *HighestCostAgents*. In April, the difference in cost between the *LowestCost* and *HighestComfort* agents is small, around $\$2.50$, due to the mild climate in Austin in April. As a result, the *BalancingAgent* does not have much opportunity to provide a large reduction in cost from the *HighestComfortAgent*, but it is still able to come close to meeting the budgets while keeping the percentage of minutes outside the preferred comfort range well under the 94.39% that the *LowestCostAgent* causes. The results for the August simulations show that the *BalancingAgent* is able to meet the given budget, while providing the user with more comfort than the *LowestCost* agent. For all simulations, when we compare the *BalancingAgent's* performance given different budgets in the same month and with the same temperature preferences, we find that the user spends more time outside the preferred temperature range with a lower budget. This is intuitive, because in order to spend less money, the HVAC must run less frequently. However, the *BalancingAgent* is still able to provide at least 40% more minutes of the month in the preferred comfort range as compared to the *LowestCost* agent.

Overall, we find that our *BalancingAgent* can provide cost savings within the range of the $\$6$ difference between the lowest cost and maximum comfort settings, in both winter and summer months, and it can do so while preserving some of the user's desired comfort. The opportunity for cost savings are diminished during the mild month of April. The potential for cost savings is limited by user preferences: a greater range between preferred and absolute settings will result in a greater opportunity for cost savings, at the expense less comfort.

IV. Conclusions and Future Work

We have presented an agent to control a residential thermostat while balancing user budget and comfort preferences. While much prior work has been devoted to taking advantage of user schedules to reduce heating and cooling costs when occupants are away, our approach focuses instead on users' common motivations to reduce the monetary cost of their heating and cooling behavior. Specifically, BALANCE is useful to those with a strong desire to keep their energy bill below a certain amount and who are willing to relax their ideal comfort requirements. We demonstrated that our simple *BalancingAgent* can reliably meet a reasonable user-provided budget[7] and provide at least a 40% improvement in comfort over the lowest-cost approach, given user preferences. There are also several avenues for improving the performance of the BALANCE agent that have the potential to meet the budget with more time spent in the preferred comfort range.

We plan to develop improved agents whose thermostat control strategies are based on machine learning techniques such as reinforcement learning and include additional contextual information such as weather forecasts and user behavior. We will use the weather forecast when determining the budget for each time step and day: if the agent knows a mild day will be followed by a hot day, it can "reserve" a higher budget for the upcoming hotter day. For more aggressive budgeting, we can incorporate the tiered pricing structures used by many utilities. We can also incorporate user schedules to reduce heating and cooling when homes are unoccupied. This is a nontrivial task for occupants who do not have regular schedules that can be programmed into the thermostat, as the system must have a way to determine when the home is unoccupied. We can look to previous work (e.g. [6]) for approaches to this problem.

Finally, we intend to implement our thermostat control system in a real world setting. A real world setting would allow us to incorporate additional learning techniques for user preferences, by monitoring manual adjustments to the thermostat. For example, if the user provides an absolute high temperature of $80°F$ but in practice frequently manually reduces the setpoint from $80°F$ to $77°F$, the agent could learn that the user's true maximum temperature should be $77°F$.

[7] The budget must be between the costs of keeping the thermostat setting at the preferred temperatures all month and keeping the thermostat at the absolute temperatures all month, if the agent is to stay within these temperature ranges.

V. Acknowledgments

This work is supported by NSF IGERT, #0966298. Any opinion, findings, and conclusions or recommendations are those of the authors and do not necessarily reflect NSF's views.

References

[1] T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry, "How people use thermostats in homes: A review," *Building and Environment*, vol. 46, no. 12, pp. 2529–2541, December 2011.

[2] C. Fischer, "Feedback on household electricity consumption: a tool for saving energy?" *Energy Efficiency*, pp. 79–104, May 2008.

[3] K. Ehrhardt-Martinez, K. Donnelly, and J. Laitner, "Advanced metering initiatives and residential feedback programs: A meta-review for household electricity-saving opportunities," American Council for an Energy-Efficient Economy, White paper, June 2010.

[4] D. Urieli and P. Stone, "A learning agent for heat-pump thermostat control," in *Proc. 12th Int. Conf. Autonomas Agents and Multiagent Systems*, May 2013, pp. 1093–1100.

[5] W. J. Cole, K. M. Powell, E. T. Hale, , and T. F. Edgar, "Reduced-order residential home modeling for model predictive control," *Energy and Buildings*, in press 2014.

[6] C. Koehler, B. Ziebart, J. Mankoff, and A. Dey, "Therml: occupancy prediction for thermostat control," in *Proc. 2013 ACM Int. Joint Conf. Pervasive and Ubiquitous Computing*, 2013, pp. 103–112.

[7] M. Gupta, S. Intille, and K. Larson, "Adding gps-control to traditional thermostats: An exploration of potential energy savings and design challenges," in *Proc. 7th Int. Conf. Pervasive Computing*, 2009, pp. 95–114.

[8] S. Lee, Y. Chon, Y. Kim, R. Ha, and H. Cha, "Occupancy prediction algorithms for thermostat control systems using mobile devices," *IEEE Transactions on Smart Grid*, vol. 4, pp. 1332–1340, March 2013.

[9] J. Scott, A. B. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar, "Preheat: controlling home heating using occupancy prediction," in *Proc. 13th Int. Conf. Ubiquitous Computing*, September 2011, pp. 281–290.

[10] M. Shann and S. Seuken, "An active learning approach to home heating in the smart grid," in *Proc. 23rd Int. Joint Conf. Artificial Intelligence*, August 2013, pp. 2892–2899.

[11] M. Power, "Low-income consumers' energy bills and their impact in 2006," Economic Opportunity Studies, White paper, October 2005.

[12] A. Meier, C. Aragon, B. Hurwitz, D. Mujumdar, T. Peffer, D. Perry, and M. Pritoni, "How people actually use thermostats," in *Proc. ACEEE Summer Study on Energy Efficiency in Buildings*, 2010.

[13] S. Chava, R. Ennaji, J. Chen, and L. Subramanian, "Cost-aware mobile web browsing," *IEEE Pervasive Computing*, vol. 11, no. 3, pp. 34–42, March 2012.

[14] D. Chassin, K. Schneider, and C. Gerkensmeyer, "Gridlab-d: An open-source power systems modeling and simulation environment," in *Transmission and Distribution Conf. and Exposition*, 2008, pp. 1–5.

[15] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin, "Aggregate model for heterogeneous thermostatically controlled loads with demand response," in *Power and Energy Society General Meeting*, 2012, pp. 1–8.