Thermovote: Participatory Sensing for Efficient Building HVAC Conditioning

Varick L. Erickson
Electrical Engineering and Computer Science
University of California, Merced
verickson@ucmerced.edu

Alberto E. Cerpa
Electrical Engineering and Computer Science
University of California, Merced
acerpa@ucmerced.edu

Abstract

Thermal comfort has traditionally been measured solely by temperature. While other methods such as Predicted Mean Vote (PMV) are available for measuring thermal comfort, the parameters required for an accurate value are overly complicated to obtain and require a great deal of sensory input. This paper proposes to bypass overly cumbersome or simplistic measures thermal comfort by bringing humans in the loop. By using humans as sensors, we can accurately adjust temperatures to improve occupant comfort. We show that occupants are more comfortable with a system that continually adjusts to thermal preference than a system that attempts to predict user comfort based on environmental factors. In addition, we also show that such a system is able to save 10.1% energy while improving the quality of service.

Categories and Subject Descriptors

H.1.2 [**User/Machine Systems**]: Human Factors

General Terms

Experimentation, Measurement, Performance

Keywords

HVAC conditioning, PMV, thermal comfort, phones

1 Introduction

Recently, works have shown how wireless sensor networks can be utilized to reduce the energy consumed by buildings[5][3] [9]. However, within the WSN community, little research has been done to improve the quality of service for users. Instead, the aim has been simply to maintain, or in some cases potentially decrease quality of service in order to achieve greater efficiency. While increasing efficiency is an important goal, the service that HVAC systems provide is arguably more important than reducing energy. Before we attempt to reduce energy for HVAC systems, we must first ensure that the system is providing adequate service.

Buildsys'12, November 6, 2012, Toronto, ON, Canada.

Building managers typically rely on building management systems (BMS) to maintain user comfort. Managers choose a single temperature set-point for occupied periods. This set-point is typically chosen based on the criteria set by American Society of Heating, Refrigeration, Air-Conditioning Engineering (ASHRAE) Standard 55 [2]. This standard uses Predicted Mean Vote (PMV), to estimate the most comfortable temperature. PMV uses multiple parameters such as humidity, temperature, and air flow to estimate how warm or cold occupants feel on a discrete scale from -3 to 3. Positive values indicates occupants are warm and negative values indicates occupants are cold. A PMV of 0 indicates occupants are comfortable. Temperature set-points are chosen by assuming fixed values of most parameters and then solving for a temperature that will give a PMV of 0.

Estimating PMV is inherently prone to error. Often, many of the values for parameters such as clothing coefficients or activity levels are given fixed values based on tables supplied by bodies such as ASHRAE. The clothing coefficient has been shown to differ by up to 20% depending on the table and method [4] [2]. Other parameters such as mean radiant temperature is currently not measured by most BMS systems and is again often estimated [13]. While it is possible that PMV estimates could be improved by attempting to measure the parameters such as occupant activity, perceived air-flow, and radiant heat for each space, this would require additional cost in terms of installation and development of specialized sensors. Even if perfect measurement is possible, the differences among individual preferences make error in the PMV estimate inevitable.

Rather than develop an entire testbed used to sensing parameters related to thermal comfort, we propose a more direct means of measuring user thermal comfort; ask the users. For this application, humans are the best sensors available. As data muling is to using humans to transport of data, participatory sensing is to using humans to sense data. We propose leveraging existing wireless infrastructure such as cell phones and use humans as sensors. By controlling room temperatures directly using user feedback, we are more likely to increase the overall thermal comfort. This is already done in private offices where individual thermostats are available. However, for rooms with multiple users, it is not straightforward to regulate temperatures from a single thermostat, such that the comfort of the majority is maximized.

In this paper we contribute the following:

- We demonstrate a learned and real-time method of utilizing user data to remove error when calculating PMV and show that both are more effective than current standard approaches of estimating PMV.
- We tested our strategies over long periods (up to 5 months) and show that temperature control based on actual feedback is substantially more effective than baseline approaches that rely solely on thermal comfort models. We show 67% and 100% satisfaction rates for our learned and real-time control strategies respectively.
- We show that improving comfort can improve energy efficiency. We found that our real-time system was able to save 10.1% energy over standard baseline strategy.

2 Related Work

While there have been sensors and sensor networks developed for measuring thermal comfort, there has been no work that has actuated temperatures using participatory sensing. Within the building community, multiple works have been done to develop different models of thermal comfort [10] and study the effects of temperature on thermal comfort [14, 12, 8]. Psychological studies have also shown that direct control can lead to greater satisfaction [11], which is an important factor for building management.

Authors of [15] developed a comfort sensor that measures air temperature, air motion, mean radiant temperature, and humidity. Input from these sensors are then used with PMV or standard effective temperature (SET*) models. They examine how their sensor choice affects the estimates of PMV as compared with ideal sensors. The main drawback is that there is no verification that compares the sensed comfort with the actual comfort. For factors such as air speed, it is unclear how many of these sensors are required for adequate estimates of thermal comfort. As occupants provide comfort level directly, our approach guarantees appropriate coverage and accurate data. Their work also does not discuss how this information can be used for temperature actuation.

In the paper [13], the authors propose an architecture for thermal comfort. They utilize SunSPOT nodes to measure air temperature. They assume mean radiant temperature is equal to dry bulb temperature and for the remaining parameters, they assume fixed values. Sensors are distributed in a room and used to estimate a PMV value at each location. These values are used to extrapolated PMV values at unknown locations within the room. They do not address how much coverage is required or how sensor coverage could potentially affect the error derived from the interpolated values. They do not compare their results with actual occupant comfort and do not discuss how these estimates are integrated with temperature control. Again, our participatory approach bypasses these issues by eliminating measurement error and obtaining measurement at the precise user location.

Authors of [7] developed a participatory sensing cell phone application for measuring temperature, lighting and air quality. For measuring temperature, they used a 5 point scale rather than a 7 point scale. They tested their application over the course of 10 days in 8 different rooms and collected 200 data points; they do not mention how many occupants participated. They compare measured temperature reading with user data and the analyze differences. They do not use this data to actuate temperature changes as we do. Our work discusses how differences from the measured PMV and sensed data can be used for correction of temperatures. Our deployment has run for a substantially longer period (5 months). Since user feedback directly impacts their environment, the incentive to participate provides us with more data. As illustration, *a single room* of our deployment over a 10 day period had 362 data points whereas their 8 room deployment only yielded 200 data points.

3 Thermovote Overview

Our deployment uses the BACnet based HVAC system installed in our building. BACnet is a client-server based communication protocol designed for the building automation and control networks [6]. The system is administrated using WebCtrl, which is a web interface developed by Automated Logic that sends BACnet commands over the network. Buildings are typically divided into zones. Each of these zones contains a collection of BACnet points where each BACnet point represents an interface to a particular sensor or point of control. For example, a room will often have a BACnet point for measuring zone temperature and another point for changing the temperature set-point. In order to use these points, we "read" or "write" commands to a point. One of the features of WebCtrl is an interface that allows the use of SOAP, an xml based protocol, to create commands to read and write BACNet points within WebCtrl. This feature is used to change temperature set-points and to monitor VAV airflow and discharge temperatures.

3.1 Design Considerations

Deciding what is a comfortable temperature is difficult. Occupants often have conflicting perceptions of what constitutes a comfortable temperature. We asked users in a survey what temperature they considered comfortable. We then compared this temperature to the actual comfortable temperature they were experiencing, we found a root mean squared error of 3.8 F^o. We decided to use the same 7 point scale as specified in ASHRAE standard 55. This scale uses values from -3 (cold) to 3 (hot) for indicating comfort. While users may not agree on a numerical temperature, they are more likely to agree on whether they feel warm or cold.

The next consideration is how to use this information. If a person states they are warm, then how much should the temperature be changed? We decided to indirectly use the Fanger's PMV model (ASHRAE 55), which is one of the most widely used and studied models for thermal comfort. Typically these models are used to predict user comfort level. In our case, we know the actual comfort level. We can thus work backwards to solve for the parameters that will give a more comfortable temperature given user feedback. This is discussed in more detail in Section 4.

Lastly, we need to address how to handle diverging opinions. What happens if one person feels cold while another feels hot? To handle this issue, we use a voting based scheme. Users are allowed one vote valued from -3 (cold) to 3 (warm) for each voting period. After each voting period, votes can be then tallied and aggregated.

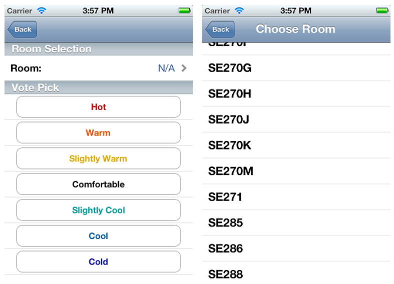

Figure 1. Screenshot of the iPhone application.

3.2 Participatory Sensors

For our application, users provide one of the following "readings"; hot, warm, slightly warm, neutral, slightly cool, cool, and cold. Users are able to provide participatory readings by using a cell phone application (IPhone and Android) or by a website. Figure 1 shows each of these input methods. Security is an issue as false feedback can negatively impact room temperatures. For the website and cell phone applications, security is achieved by authenticating through the school central authentication service (CAS). To prevent people for voting in rooms they do not occupy, we require users to submit a room request, which we manually verify to ensure they occupy the space. Once a user room request is approved, the user may provide thermal feedback for that room. Rooms are stored so users do not have to repeatedly enter their choice, reducing the burden of providing data.

As with any participatory sensing application, privacy is an issue that should be considered. Since thermal comfort level is not considered sensitive, most occupants feel comfortable providing this information. Occupants would have to report this information regardless if they were to request temperature change from management. Erring on the side of caution, we chose not to include any identifiers when storing user feedback in our database.

4 Temperature Control

This section examines a real-time and learned strategy for temperature control. We start with a discussion of how baseline temperatures are typically established and how PMV is utilized. We then show a method of using participatory sensing to achieve temperature control in real-time.

4.1 Baseline Control

Determining a set-points for temperatures is often an imprecise exercise. ASHRAE 55 specifies guidelines for creating comfortable temperatures through the PMV metric. To calculate the PMV, the following parameters are required: metabolism (M), external work (W),mean radiant temperature (T_{rad}), air temperature (T_{air}), relative Humidity (h), partial water vapor pressure (P), clothing coefficient (f_{cl}), air velocity (v_{air}), outer clothing temperature T_{cl}. These parameters are used in the following to calculate PMV:

$$
\begin{aligned}
&P\hat{M}V(M,w,T_{rad},T_{air},h,C,v_{air}) = (0.028+0.303e^{-0.036M})\{(M-W)\\
&-3.05-3(5.733-6.99\times10^{-4}(M-W)-P)-0.42((M-W)-58.15)\\
&-0.017M(5.867-P_a)-0.0014M(34-t_a)-3.96\times10^{-8}f_{cl}((T_{cl}+273)^4\\
&-(T_{rad}+273)^4)-f_{cl}\times(T_{cl}-T_{air})\} \qquad (1)
\end{aligned}
$$

Parameter	High Solar Gain	Low Solar Gain
Metabolism	70 W/m^2	70 W/m^2
MRT	75 F^o	72 F^o
Relative Humidity	30%	30%
Clothing	1.0	1.0
Air Velocity	0.1 m/s	0.1 m/s

Table 1. Parameters used for initial PMV estimate and temperature set-points.

From a practical standpoint, only air temperature and relative humidity are commonly sensed by a BMS. Air velocity can be measured, but often differs depending on occupant location making it difficult to determine the average airflow for occupants. Mean radiant temperature (MRT) is the average weighted temperatures of the surroundings. Solar gain heating the surrounding surfaces of a room will affect the MRT. MRT can be measured using a black globe thermometer, but these sensor are typically not included in BMS deployments. The remaining parameters such as metabolism and the clothing coefficient are typically estimated.

The initial heating set-points for our building during the winter months was 74 F^o/23.3 C^o for rooms with little solar gain. Rooms with more exposure to solar gain had heating set-points of 72 F^o/22.2 C^o. These set-points were determined by estimating values for all the parameters except air temperature (see Table 4.1). One drawback to this strategy is that since many of parameters are fixed, the error of the PMV will vary depending on changing environmental factors. For example, the PMV may be accurate in the morning for a particular room when there is no sun. However, if the room receives more sun in the afternoon, then the error of the PMV estimate will increase. The following strategy is designed to dynamically correct these PMV errors.

4.2 Learned Control Schedule

Since many of the parameters of PMV system are estimated, these estimates can differ from actual comfort. If we collect data from users, we can compare the estimated PMV with the actual comfort and create a temperature correction factor. For example, if the estimated PMV is 0, but the measured value from occupants is -1, then can increase temperatures to offset this PMV error of -1. This can be done as follows,

$$P\hat{M}V(M,w,T_{rad},T_{air}+T_{offset},h,C,v_{air}) - AMV = 0 \qquad (2)$$

where T_{offset} is a temperature offset that will correct the initial estimate as compared with the actual mean vote AMV.

The intuition behind the learned control strategy is that it uses collected user feedback to learn the temperature correction offsets for different parts of the day. For example, a room that feels warmer in the afternoon would have an temperature adjustment for that period. We can adjust temperatures dynamically throughout the day according to these offsets and help correct for changes in the environment. Algorithm 1 shows the algorithm for the learned control schedule. Every n minutes (10 minutes), we examine the previous user votes for this particular window of time and find the actual mean vote. We then find the T_{offset} to correct for the initial PMV estimate. We then compare the current air temperature T_{air} with the current set-point $T_{setpoint}$. If the difference between T_{air} and $T_{setpoint}$ is less than a certain threshold, then we know that the air temperature is currently at the specified set-point and we then change the set-point by T_{offset}. We do not change the set-point if the difference is above the threshold; this is because the system is in the process of bringing the room to the specified set-point. This threshold should be set to the amplitude of temperature change caused by the PID controller. To prevent voting bias, for each user, we only count the most recent vote for each 10 minute period.

While the learned control strategy is an improvement over the baseline strategy, there are still limitations with this approach. The main limitation of this approach is that the corrections are based on the time when user data was collected. For example, if the data was collected in the winter, then applying the same corrections during the summer may not be valid. This seasonal error could be minimized by collecting data at different times of the applying corrections accordingly. However, there are still situations such as heat waves, where the historical corrective offsets may be inaccurate. The main strength of this approach is that it does not require an advanced BMS. Most modern thermostats are capable of creating daily temperature schedules. One could pre-compute the corrected set-points and enter the temperatures into the system.

Algorithm 1 Learned Control Schedule

$\hat{PMV} \leftarrow$ See Equation 1
$AMV(i,j) \leftarrow$ Actual Mean Vote from previous data from time i to j
$T_{setpoint} \leftarrow$ Current temperature set-point
$t_i \leftarrow$ Time at instance i
$thresh \leftarrow$ Threshold to consider $T_{air} \approx T_{setpoint}$

for Every n minutes **do**
  $t_{current} \leftarrow$ Current time

  Solve for T_{offset}:
  $\hat{PMV}(M, w, T_{rad}, T_{air} + T_{offset}, h, C, v_{air}) + AMV(t_c, t_{c+n}) = 0$

  if $|T_{setpoint} - T_{air}| < thresh$ **then**
    $T_{setpoint} = T_{air} + T_{offset}$
  end if
end for

4.3 Real-Time Control

The real-time control strategy performs the same correction as the learned strategy except the corrections are done in real-time. Every n minutes (10 minutes), we examine the previous user votes for this particular window of time and find the actual mean vote and find the T_{offset} to correct for the initial PMV estimate. We then compare the current air temperature T_{air} with the current set-point $T_{setpoint}$. If the difference between T_{air} and $T_{setpoint}$ is less than a certain threshold, then we know that the air temperature is currently at the specified set-point and we then change the set-point by T_{offset}. We do not change the set-point if the difference is above the threshold; this is because the system is in the process of bringing the room to the specified set-point. This threshold should be set to the amplitude of temperature change of the PID controller. To prevent voting bias, for each user, we only count the most recent vote for each 10 minute period.

5 User Studies

For our study, we recruited 39 participants to provide feedback. The building was constructed within the past 6

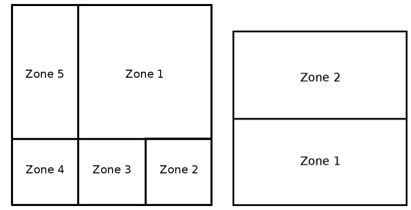

(a) Office (b) Lab

Figure 2. Layout of Office and Lab.

years and is LEED Gold certified. We ran the baseline, learned, and real-time system in an administrative office (5 zones, Figure 2(a)) and a graduate lab (2 zones, Figure 2(a)). There were a total of 39 participants (all the personnel in these areas) and the study was conducted over 5 weeks. The study was conducted in the winter where the system was always in heating mode.

***Study 1, Baseline Evaluation*:** The first week, was used to determine the actual performance of the static schedule. During this period, we asked occupants to log their comfort level. The purpose of this evaluation was collect data for creating a learned control schedule and to examine baseline performance. At the end of the week, we sent out a survey to gather information regarding their past experiences regarding thermal comfort within the building and overall satisfaction.

***Study 2, Learned Control Schedule*:** User feedback from the first week allowed us to correct our initial estimate for PMV. We then created a new dynamic temperature schedule using the method described in Section 4.2. Occupants again were asked to log their their comfort level so we could compare the corrected PMV with their actual comfort level. At the end of the week, we sent another survey asking about temperatures during the week and their overall satisfaction with the new temperature schedule.

***Study 3, Real-Time Control*:** 3 weeks were used to test the real-time control strategy described in Section 4.3. At the end of this period we again sent another survey regarding their satisfaction with the system.

***Study 4, Long Term Real-Time Control*:** 5 months were used to test the real-time control strategy described in Section 4.3. At the end of this period we again sent another survey regarding their satisfaction with the system.

6 Results

For our analysis, we examine three different aspects for each strategy. We start with the performance of each strategy with respect to thermal comfort. We then examine maintenance issues for this type of system. Finally we examine how the system changes the energy consumption.

6.1 Thermal Comfort

We start by examining the baseline comfort level of occupants. Figure 3 shows the votes for several rooms over a 5 week period as a function of time. During the baseline period, most occupants felt Cool to Cold as can be seen from the votes in Figure 3. Figure 4(a) shows the initial PMV estimate based on the fixed parameters described in Table 4.1 for the high solar gain. We see that this initial estimate is close to zero (Neutral) most of the time, however, since the BMS is not always able to meet the target set-point, there

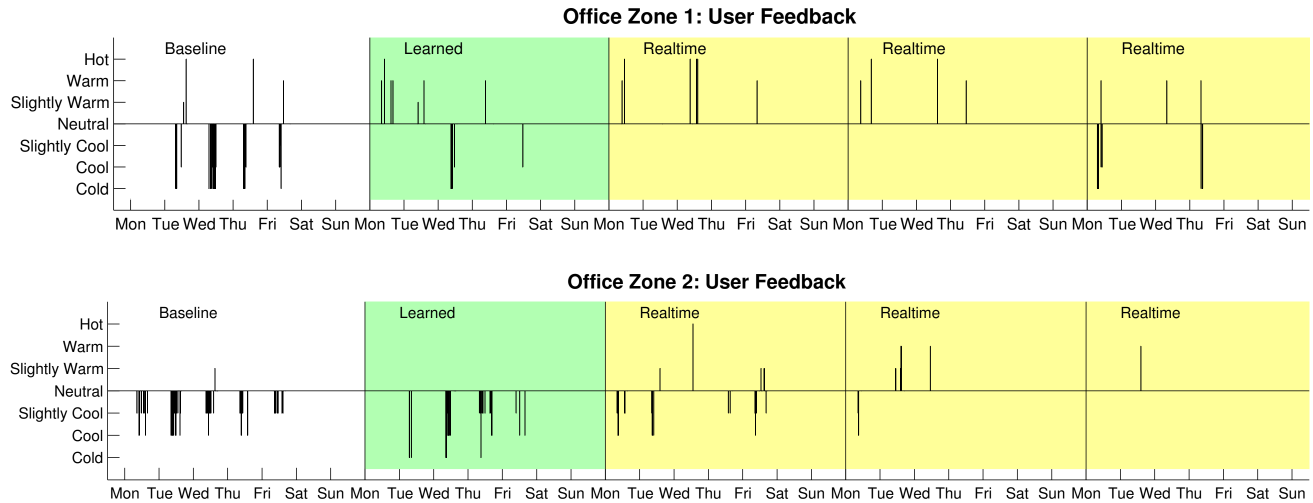

Figure 3. Study voting patterns for the different control strategies for two representative zones

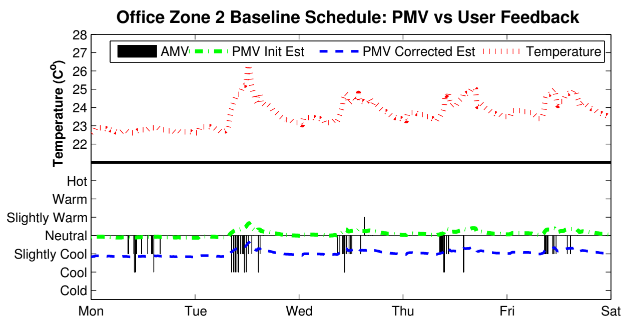

(a) Baseline

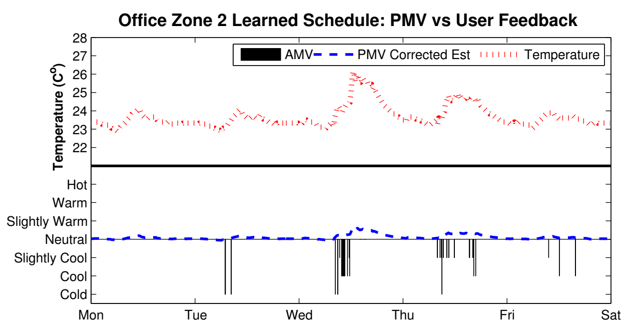

(b) Learned

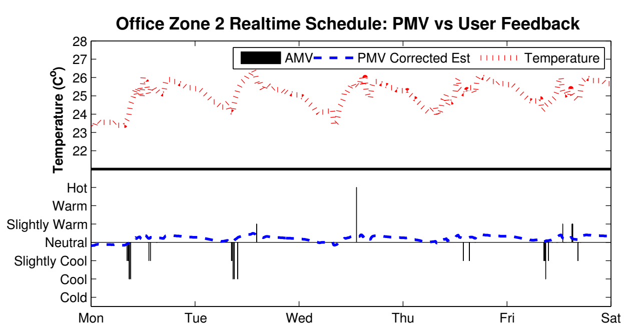

(c) Realtime

Figure 4. The initial PMV estimate, corrected PMV, votes, and temperature for Office Zone 2.

(a) Thermal Comfort	Baseline	Learned	Real-time
Cold	29%	27%	0%
Cool	29%	9%	0%
Slightly Cool	18%	9%	22%
Neutral	0%	19%	67%
Slightly Warm	12%	27%	11%
Warm	12%	9%	0%
Hot	0%	0%	0%

(a) Satisfaction	Baseline	Learned	Real-time
Dissatisfied	33%	8%	0%
Somewhat Dissatisfied	42%	17%	0%
Neutral	0%	8%	0%
Somewhat Satisfied	17%	50%	77%
Satisfied	8%	17%	23%

Table 2. Office thermal comfort and satisfaction.

is some variability in the PMV estimates and it fluctuates around zero. Nevertheless, a value so close to zero should indicate that the occupants are comfortable. However, if we examine the votes during this period, we see that occupants actually voted that they were cool to cold. The survey results also reflect this finding. The summed percentage cold to cool for baseline (Table 2(a) column 1) is 76%. No occupant indicated that they felt comfortable. The satisfaction survey reflected this showing that 75% of occupants were not satisfied with the conditioning of the room. This indicates that there is error in the PMV estimate and a positive temperature offset is required to correct the PMV estimate. We estimated the AMV by averaging the votes over 10 minute periods. The total we average over is the number of unique occupants who have voted for the current day. If occupants do not vote during a period, then it is assumed they voted neutral –this is because occupants seldom ever report they are comfortable, but are eager to report when they are not–. Once calculated, we can then use the average vote to correct the PMV estimate. The PMV Corrected Estimate line in Figure 4(a) shows the PMV after correcting for the error.

The second week where the learned strategy was applied showed improved results. Office Zone 2 showed AMV values closer to zero (Figure 3). The decrease in votes indicates the PMV estimate learned from the prior week is closer and that temperatures are more comfortable. Figure 4(b) shows

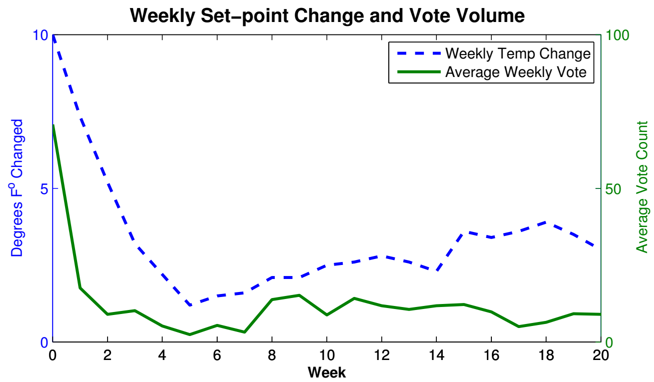

Figure 5. Weekly vote volume and total temperature change over 5 months.

the corrected PMV estimate for Office Zone 2, which is close to zero. In Office Zone 1, however, while the total number of votes decreased, there were more votes that were slightly warm to hot (Figure 3). This is likely due to the temperature conditions of the baseline week. If temperatures of the baseline week are significantly different than temperatures during a given week, then the PMV correction may not be valid. Overall, however, the surveys showed occupants to be more comfortable. Table 2(a) column 1 shows 19% of occupants were comfortable as opposed to 0% for baseline. We also see that less people indicated that they felt slightly cool to cold; only 45% percent felt cold to slightly cool versus the 76% found for baseline. Occupants were also more satisfied with the learned temperature schedule; 67% total were satisfied or somewhat satisfied (Table 2(b), column 2) whereas only 25% were satisfied or somewhat satisfied with the baseline (Table 2(b), column 1).

The real-time strategy showed significantly improved thermal comfort and satisfaction for the three week short term study. Figure 3 shows the frequency of the votes decreased 67% from the learned schedule to the first week of the real-time study. Figure 4(c) shows the corrected PMV estimate is close to 0. This is expected since occupants are correcting the PMV estimate in real-time. When conditions change, they are able to account for the change through voting. Again, this is also reflected in the surveys. 67% felt comfortable, 22% felt slightly cool, and 11% felt warm (Table 2(a), column 3). Table 2(b) column 3) shows that 77% were satisfied and 23% were somewhat satisfied; 0% of occupants were dissatisfied or somewhat dissatisfied.

In addition to the three weeks of real-time results provided by the office areas, we also have 5 months of extended real-time strategy results from Lab 1. Here we also saw significant improvements. Table 3(b) column 1 shows that 100% were not satisfied with the baseline service. After using the real-time strategy, the results showed that 80% were satisfied with the real-time system. Baseline temperatures for Lab 1 where split; almost half of the lab was cold (46%) while the other half (54%) was warm (Table 3(a); 0% felt comfortable. For the real-time strategy, 53% felt comfortable. In general, people who did not feel comfortable perceived colder temperatures (25% felt slightly cold to cold); 13% felt slightly cool, 6% cool, and 6% cold. 19%

(a) Thermal Comfort	Baseline	Real-time
Cold	23%	19%
Cool	23%	11%
Slightly Cool	0%	16%
Neutral	0%	38%
Slightly Warm	15%	11%
Warm	14%	5%
Hot	25%	0%
(a) Satisfaction	Baseline	Real-time
Dissatisfied	46%	0%
Somewhat Dissatisfied	54%	7%
Neutral	0%	13%
Somewhat Satisfied	0%	27%
Satisfied	0%	53%

Table 3. Lab 1 thermal comfort and satisfaction.

of the people felt slightly warm to hot. Interestingly, despite 47% of the people not indicating neutral temperature, 80% of the occupants were satisfied with the system indicating that absolute comfort is not required for user satisfaction. Figure 5 shows the weekly vote volume variation over the past 5 months (right y axis). For this graph, we only include the actual votes and do not consider implicit neutral votes. Initially, we see quite a few people voting at the beginning. We can see that after the first month, the mean and variance of the votes remain stationary. This is because as temperatures become more comfortable, the number of votes decreases and then eventually stabilizes. This is also reflected in the total temperature change experienced each week (Figure 5, left y axis). This figure shows the total amount of temperature set-point change experienced for each week. At the beginning, the set-points were changed a total of approximately 10 degrees, which corresponds with the initial high vote volume. As the weeks go on, the amount the set-point changes each week gets reduced; between weeks 5 and 14, the total weekly change is only between 1.2-2.5 F°. The weekly change increases somewhat (3-3.6 F°) during weeks 15 to 20. This is due to Spring/Summer season change. Overall, if we do not include the first month, we see that the total amount of set-point changes is between 1.2-3.9F°. These long term results show that even after the initial "novelty" period, occupants still use the system consistently.

These results also provide some interesting insight regarding thermal comfort level for majority. For the real-time system, we initially see many votes. Over time, the number of votes decreases and then the rate remains fairly constant. Once the temperature converges to a comfortable level, it tends to remain static. This is particularly true for areas with fewer occupants and can be seen in Figure 3 for Zone 1 and Zone 2. For larger areas such as the lab, the vote rate tends to be larger as it is more likely to have greater variety for thermal preference; however, the temperature set-point still remains fairly static.

This raises the question of perceived thermal comfort consistency. At the individual level, occupants are 100% accurate when determining if they are comfortable; only their opinion matters. However, issues of temperature control can arise if users have diverging opinions for room temperatures. Figure 6 shows the variance among user votes during the baseline period. For Office Zones 3 and 4, votes are ex-

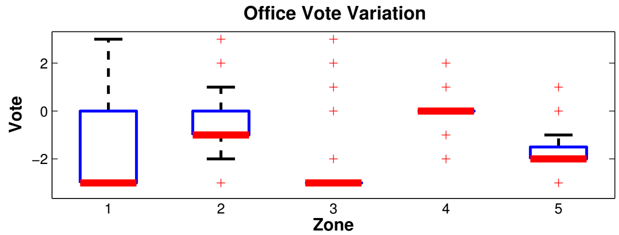

Figure 6. Variation of votes for zones occupied by more than one occupant. The dark thick line (red) shows the median vote. The top and bottom of the box are the 25th and 75th quantiles. The "whiskers" show extreme data not considered outliers. Outliers are plotted as "+".

tremely consistent among occupants. For these zones we only see a few outliers. Office Zone 1 shows the largest variation. Most occupants voted "Cold" (-3), but enough "Hot" votes were recorded during this period that these readings are not considered outliers. One possibility is that this difference could be amplified by competing thermal preference. The minority of occupants not wishing the temperatures to increase may attempt to offset the majority's vote by voting for the extreme opposite. We also see greater variation for Office Zones 2 and 5, but see the differences are smaller.

The convergence of temperatures over time also suggests that many people hold a similar feeling for what is comfortable. However, we suspect that the scale between neutral to hot and neutral to cold differs significantly among occupants. What one occupant might perceive as slightly warm another may perceive as hot. Despite this difference they are still likely to have the same perception of what is comfortable. Since users can feel the results from their own input, they can re-correct the temperatures until they feel comfortable; they essentially function as part of a proportional integral derivative controller eliminating the input error.

6.2 Maintenance and Management

In addition to thermal comfort, another important factor to consider is system management. For building management, reducing the number of complaints is of critical importance. In our surveys, we also included questions regarding users' experience interacting with facilities management.

During the learned and real-time study period, users only contacted us twice. The first issue regarded extreme temperatures for three areas connected to a shared VAV unit. In this configuration, a single VAV conditions the air for multiple rooms. The proportion of air each room receives from the VAV is controlled by manually adjusting each room damper. In this case, one room was receiving approximately 200 CFM the air and each of the remaining rooms were only receiving 50 CFM. Once the cause was determined facilities was informed of the issue. The other was when a network outage caused our vote system to go off-line. Based on the survey results, 100% of occupants were either satisfied or somewhat satisfied with our resolution.

In addition to issues reported by users, we were also able to correct other issues by examining votes. During one afternoon, we noticed a sudden increase of votes indicating overly warm temperatures in the lab area. When we investigated further, we discovered one of the VAV's connected

(a) Contact for past year

Did Not Contact	37%
Unable to Contact	6%
Did Contact	56%

(b) Satisfaction with resolutions

Dissatisfied	33%
Somewhat Dissatisfied	2%
Neutral	0%
Somewhat Satisfied	33%
Satisfied	11%

Table 4. User interaction with facilities management before Thermovote.

to the BMS was locked in a default state. This default state used overly warm temperature set-points and prevented set-point change commands from being accepted. Over the first few weeks, this occurred on multiple occasions. Once aware of this state, we began including checks to ensure that out set-point changes were being accepted, and configured the system to email us when the state occurred. Without the vote feedback, it would have been difficult for facilities to know the problem existed unless they happened to examine that specific VAV or if a user contacted facilities.

Prior to Thermovote, 56% of occupants contacted facilities regarding room temperature during the past year (Figure 4). Of these occupants, 44% were satisfied or somewhat satisfied. Given that 76% of occupants were dissatisfied with thermal comfort in the office (Table 2(b)) and only 56% actually contacted facilities, we estimate 24% choose not to contact facilities. Since that majority of occupants who contacted facilities were not satisfied with the results, it is possible that these 24% occupants view contacting facilities as ineffective method of adjusting temperature. This is not entirely surprising. Even if facilities is attentive, responding to each complaint, if they do not collectively incorporate all feedback from the users, it is unlikely set-points will converge to a comfortable temperature.

6.3 Energy Consumption

In this section we examine how the different strategies affects energy consumption of the HVAC system. Rooms are conditioned primarily using variable air volume (VAV) units. The main air handler unit drives air to each of the VAV units. Air moves through a water coil in each of these units, heating or cooling the air. To measure energy we use the following equation: $Q = mC_{air}(T_{in} - T_{out})$. Q is energy transferred from the coil to the air, m is the total air mass passing through the coil, C_{air} is the specific heat of air, T_{in} is air temperature of the incoming air, T_{out} is the temperature of the outgoing air after passing through the coil. These three parameters are measured by the BMS. By measuring the airflow and the supply and discharge temperatures, we are able to calculate the amount of energy consumed by each VAV. The mass m can be calculated by measuring the airflow over a period of time to determine the volume of air and then multiplying the volume by air density. Airflow, T_{in}, and T_{out} are measured measured by the BMS.

Interestingly, the real-time system showed 10.1% savings over the baseline strategy; this shows that in certain situations a real-time system could potentially save energy. To account for energy differences caused by factors such as weather and humidity, baseline energy consumption was calculated from similar days from historical data. The baseline week had an average temperature of 55.0 F^{o} with a standard

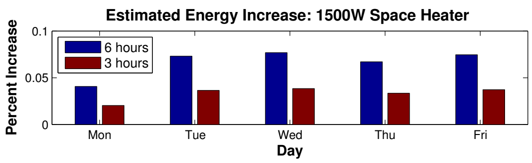

Figure 7. The energy cost of using a space heater.

deviation of 1.5 F^o. The real-time week used had an average temperature of 54.4 F^o with standard deviation 1.9 F^o. While the temperatures in the office zone increased on average 2.1 F^o, temperatures in the lab zones decrease on average 3.3 F^o. Room temperatures tended to be similar causing less "thermal competition" between rooms and increasing overall efficiency. For example, a room conditioned to 72 F^o next to room conditioned to 76 F^o will cause additional expenditure of energy due to the thermal transfer between the walls.

Another factor to consider is personal conditioning devices such as as space heaters. During a walk through of the area, we found 4 heaters ranging from 1000W to 1500W. We do not have data regarding the actual usage of the heaters, but provide a rough estimate of their potential energy impact. We used the energy consumption data from [1] for representative 1500W space heater, which uses 950W to maintain 70 F^o. 10 hours, then 15KW more energy for each Figure 7 shows the energy increase caused if the heater is different amounts of time per day (6 hours, 3 hours) over the course of a typical week. This is a 7% increase in energy for the week for approximately 33% of the zones. If extrapolated over the entire building, space heaters potentially consume a significant amount of energy that could be saved through proper conditioning.

7 Discussion

In this section, we discuss some of our experiences and observations while utilizing participatory sensing. Unlike traditional sensing, when using participatory sensing, one is at the mercy of the users. One concern for this application was participation. Would users be willing to provide feedback for an entire week? Would enough data provided be sufficient to estimate the AMV in order to condition effectively?

We found participants to be very enthusiastic about providing votes. Indeed, some participants actually voted every 15 minutes for the entire first week; occupants wanted to ensure their thermal welfare. Unlike some participatory sensing applications, for this particular application, users have a vested interest to provide feedback. We believe that this is an important factor for the success of a participatory sensing application. This interest however can be a double edged sword. There is evidence that some people can provide spurious feedback. By over-inflating estimates of their comfort level, users hope to give their preference more weight and thus better comfort. While in this application, this did not seem to affect overall user satisfaction, this must be taken into consideration when attempting other participatory sensing applications. Human perception of control must not be overlooked. Studies have shown that users who perceive they have control are more likely to be satisfied [11]. Our results seem to confirm that some degree of control greatly increases satisfaction and perceived comfort.

For our application, one method of reducing this bias would be to change the voting scale from a 7 point scale to a 3 point scale. This would only allow users to choose "hot", "neutral",and "cold" removing possibility of bias. However, this would also remove useful information. While this approach would remove the over weighting of outliers, it would also remove the comfort level of the majority; we would no longer be able to tell if the majority is slightly warm or hot.

8 Conclusion

In this paper, we developed Thermovote, a temperature control system that utilizes participatory sensing in order to actuate temperature change. We developed IPhone and Android applications along with a website that allows users to provide feedback regarding their thermal comfort and show a real-time method of using user data for controlling the temperatures of rooms. This is achieved using a method of correcting PMV estimates using AMV in order to determine temperature changes. We tested our real-time system over a period of 5 months. We also ran three studies over a period of 5 weeks testing the learned and real-time strategies. We showed 100% satisfaction for the real-time strategy whereas only 25% were satisfied with the baseline strategy. In addition, we show that adjusting thermal comfort can actually increase efficiency. For our deployment we show 10.1% energy savings over the baseline strategy.

9 References

[1] Appliance data energy use. `http://minotaur.lbl.gov/aeud/appliances/space-heater/delonghi/1193/`.

[2] ASHRAE standard 55: Thermal environmental conditions for human occupancy. ASHRAE,Inc., 2004.

[3] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng. Duty-cycling buildings aggressively: The next frontier in HVAC control. In *IPSN'11*.

[4] G. Brager, M. Fountain, C. Benton, E. A. Arens, and F. Bauman. A comparison of methods for assessing thermal sensation and acceptability in the field. *Energy and Buildings*, 1993.

[5] V. L. Erickson, M. Á. Carreira-Perpiñán, and A. E. Cerpa. OBSERVE: Occupancy-based system for efficient reduction of HVAC energy. In *IPSN'11*.

[6] B. A. D. C. P. for Building Automation and C. Networks. `http://www.bacnet.org/`.

[7] B.-G. B. Jazizadeh F, Kavulya G, Klein L. Continous sensing of occupant perception of indoor ambient factors. *ASCE Workshop of Computing in Civil Engineering*, 2011.

[8] P. Kunze and J. Grunewald. Suitable algorithms for practical assessment of indoor climates in hospital wards. *Symposium on Building Physics*, Sept. 2010.

[9] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. In *SenSys*, 2010.

[10] F. P. O. *Thermal comfort. Analysis and applications in environmental engineering*. Danish Technical Press, 1970.

[11] M. Paciuk. *The Role of Personal Control of the Environment in Thermal Comfort and Satisfaction at the Workplace*. Environmental Design Research Association, 1990.

[12] J. U. Pfafferott, S. Herkel, D. E. Kalz, and A. Zeuschner. Comparison of low-energy office buildings in summer using different thermal comfort criteria. *Energy & Buildings*, 2007.

[13] M. Rawi and A. Al-Anbuky. Passive house sensor networks: Human centric thermal comfort concept. *In ISSNIP*, 2009.

[14] M. Schumann, A. Burillo, and N. Wilson. Predicting the desired thermal comfort conditions for shared offices. In *ICCCBE*, 2010.

[15] G. Ye, C. Yang, Y. Chen, and Y. Li. A new approach for measuring predicted mean vote (PMV) and standard effective temperature (SET). *Building and Environment*, 2003.