

Model-Free HVAC Control Using Occupant Feedback

Sean Purdon
University of Queensland
St. Lucia, QLD, Australia
Email: sean.purdon@csiro.au

Branislav Kusy, Raja Jurdak
CSIRO ICT Centre
Pullenvale, QLD, Australia
Email: {brano.kusy, raja.jurdak}@csiro.au

Geoffrey Challen
SUNY Buffalo
Buffalo, NY, USA
Email: challen@buffalo.edu

Abstract—Optimal control of Heating, Ventilation, and Air Conditioning (HVAC) is an important step towards reducing the carbon footprint of buildings and requires balancing energy reductions and occupant comfort. Conventional thermostats for temperature set points provide a centralised single point of user input, often leading to significant thermal discomfort for occupants. We propose to instead include users in the HVAC control loop through distributed smart-phone based votes about their thermal comfort for aggregated control of HVAC. Unlike existing approaches that require in-situ sensors or build complex comfort models of individual users, we propose a model- and sensor-free HVAC control algorithm that uses simple user input (hot/cold) and adapts to changing office occupancy or ambient temperature in real time. We develop an iterative data fusion algorithm that finds optimal temperature in offices with multiple users and propose techniques that can aggressively save energy by drifting indoor temperatures towards the outdoor temperature. Our evaluation is based on empirical data collected in 12 offices over a 3-week period and shows that adaptive HVAC control can save up to 60% of energy at a relatively small increase of 0.3°C in average occupant discomfort.

I. INTRODUCTION

Reducing energy consumption of residential and commercial buildings is an important problem that has attracted the attention of several research groups in recent years [9]–[12]. Heating, ventilation, and air-conditioning (HVAC) has been singled out as the most important contributor to greenhouse gas emissions; it comprises as much as one-third of energy consumption in the United States [7]. Some studies suggest that up to 30% of this energy can be saved by duty-cycling HVAC when residents are sleeping or away [13].

As a result, the most popular approach to improving energy usage of HVAC is to aggressively duty-cycle its operation based on the occupancy of air-conditioning zones. A variety of wireless networked sensors, such as cameras, passive infrared (PIR), or door sensors have been used to estimate building occupancy. Once room or zone occupancy can be established, the HVAC system can be duty-cycled in real-time to save energy [2], [16]. More advanced approaches use multivariate Gaussian and Markov chain models [7] or hidden Markov models [13] to learn and predict occupancy patterns based on historical data to achieve further energy reduction. Overall, the studies have demonstrated dramatic reductions of up to 42% in HVAC energy usage in commercial buildings [7] and 28% in residential buildings [13]. Furthermore, users tend to change

their behavior towards sustainability if given better visibility of the impact and cost of their actions [10].

Achieving high energy savings without considering thermal comfort of users is unrealistic. A simple strategy of setting the temperature setpoint to the ambient temperature would achieve the maximum savings, but could also have negative impact on the comfort and productivity of people. Most techniques, therefore, optimize energy consumption of HVAC within the acceptable thermal comfort bounds, as defined, for example, in the American Society of Heating, Refrigerating and Air-Conditioning (ASHRAE) comfort standards [3].

Thermal comfort of individuals is a subjective metric and is commonly modeled solely through temperature. The simplistic temperature-based comfort estimation, however, may yield errors and thus more advanced models for thermal comfort have been proposed, such as Fanger’s Predicted Mean Vote (PMV) [14]. Both these models require accurate sensing of environmental parameters in each air-conditioning zone, which typically incurs high overheads for installation and maintenance of sensors. Furthermore, some thermal comfort parameters, such as clothing insulation and metabolism levels, are subjective and difficult to measure in an automated way. Commonly these parameters are estimated by experts [4] or are obtained through user feedback [17], which is expensive and not scalable. Reliable thermal comfort estimation in field deployments thus remains one of the main barriers to widespread use of advanced HVAC technology today.

In this work, we propose a model- and sensor-free approach for estimating thermal comfort of people and optimal control of HVAC by placing humans in-the-loop. We build on ideas presented in *Thermovote*, a smart-phone based participatory sensing of thermal comfort developed by Erickson et al [8]. The beauty of our approach is that it requires neither the deployment of additional environmental sensors in the building nor complex models for estimating thermal comfort of users, as it works with user comfort input directly and is capable of estimating building occupancy automatically using smart-phone data.

In contrast to other participatory sensing techniques, such as *Thermovote*, we do not attempt to model the PMV comfort metric based on user surveys. Instead, we take a simple approach and manipulate temperature directly based on user input, i.e., we increase the indoor temperature when people




Fig. 1. Model-free and sensor-free HVAC control using smart-phone feedback.




Fig. 2. User interface of our iPhone application.

are cold and decrease it when they are hot. We propose an iterative algorithm to find the optimal temperature for a group of people based on real-time user input. The model-free aspect of our algorithm allows us to use a considerably simpler user input, than the relatively cumbersome 7-point ASHRAE comfort metric adopted by the industry. In addition to optimizing the temperature to maximize user comfort, we also implement several mechanisms to reduce HVAC energy consumption. Specifically, letting the indoor temperature drift towards the ambient outdoor temperature saves a significant amount of energy at a relatively minor increase in the required user input.

We evaluate our algorithms in empirically-based simulations. We built a wireless sensor network to measure temperature, humidity, and office occupancy of a building with 3 HVAC zones and developed a PC application for surveying thermal comfort of building occupants. We also obtained access to our building management and control system (BMS), through which we control temperature in individual HVAC zones. We use empirical data gathered in this manner to evaluate the performance of different HVAC control algorithms. We demonstrate that the optimal temperature in shared offices can be found reliably based on simple multi-user input. Our algorithm adapts well to people arriving and leaving the office and to changes in the outdoor ambient temperature in different weather seasons. Compared to the existing HVAC control in our building, our basic participatory system reduces energy usage of HVAC by 10% at a slightly worse thermal comfort of people. This is despite the fact that the existing baseline HVAC control sets the temperature to an optimal value based on user feedback and our building follows a strict policy of turning off the HVAC outside office hours. Our additional energy saving mechanisms, which we refer to as *Drift*, help to decrease the energy usage of HVAC by a further 50% while requiring slightly more user votes.

II. SYSTEM ARCHITECTURE

We propose a model- and sensor-free system to control HVAC through participatory sensing of thermal comfort of users. As shown in Fig. 1, the system relies on two main sources of information: (1) the building control and management software (BMS); and (2) an application installed on users' smart phones. This application functions as an implicit occupancy sensor.

Where possible, data from the BMS (including the current indoor and outdoor temperatures) is shown to the user. However, our algorithm does not require that such information be present. It is sufficient that the BMS provide a means of changing the temperature set-point, and that some source of outdoor temperature data is available. If the BMS does not provide information on the outdoor temperature, online sources may be used.

Sensed data is processed on a centralized server which takes two main courses of action. It changes indoor temperature to satisfy thermal comfort of the occupants and generates data that is displayed on smart-phones in real time. We aim to promote behavior change towards greater sustainability, give users a sense of ownership over the HVAC control, and enable groups of people to reach a temperature compromise.

The system closes the loop between sensing and actuation on a building scale without requiring the deployment of any additional occupancy or temperature sensors. Those sensors which are already present for the BMS are harnessed to provide user feedback.

A. Smart-Phone Application

We have developed a prototype iOS application as shown in Fig. 2. The application provides an interface for submitting the current comfort vote using simple icons in the bottom of the screen. Users can vote at any time. However, as we will explain in Sec. III only the last vote of a user in a given period is counted to prevent users taking advantage of the system.

The application also plots historical data for indoor and outdoor temperatures as well as the current HVAC temperature set-point. This information provides users with feedback on their actions and gives them detailed visibility into the current state and performance of the HVAC system. The temperature data is obtained from the BMS. If outdoor temperature is not available through BMS, it can be downloaded from weather forecasting websites. This enables operation in a truly sensor-free environment.

Finally, we show a pie-chart of the thermal preferences of all other people in the office. We expect people to moderate their votes to reach an office-wide compromise on the temperature set-point.

B. Data Collection and Fusion Server

We implemented a data fusion algorithm running on a central server for estimating the average discomfort of people in the office based on the data collected from users. We work directly with user comfort data and define the optimal temperature as the one that balances the number of people

that are too hot and too cold. If the two categories become unbalanced, for example, when the temperature or occupancy changes, we iteratively adjust the HVAC temperature until a new equilibrium is found. The system has a secondary goal of reducing energy usage while staying within the acceptable thermal comfort bounds. This maximises acceptable energy savings. More details are provided in Sec. III.

Our central server has multiple functions. In addition to actuating HVAC based on the fused user input, it also compiles information that is displayed on the phones, to provide users with feedback on the group voting process and on the HVAC status. We rely on BMS to provide us with the control of indoor temperature set-points, sensing of indoor temperature, and duty cycling of HVAC in individual zones. In our experience, even the most basic HVAC control systems expose such information through a programmatic interface.

It is important to note that the use of this data does not detract from the sensor-free nature of our system. Our algorithm merely requires that the BMS provide a means of changing the temperature set-point. This is a fundamental requirement for a useful BMS. Additional information, such as indoor/outdoor temperatures, is desirable in that it enables richer user feedback, but is not required.

III. DATA FUSION STRATEGIES

In this section, we introduce a data fusion algorithm that optimizes the thermal comfort of a group of users, while simultaneously achieving energy savings. The algorithm is model-free, in that it does not estimate the temperature at which users feel comfortable. Instead, it iteratively changes the temperature to compensate for user discomfort, simply by increasing or decreasing the temperature when users become too cold or hot, respectively.

A. Thermal Comfort

The American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE) defines thermal satisfaction to depend on both physiological and psychological factors, which makes it difficult to measure through sensors [3]. A popular method to measure thermal comfort is the industry-standard Fanger's PMV [14], a measure on a 7-point scale between cold (-3) and hot (3). The method depends on six variables that include person-dependent metabolic rate and clothing level. In addition to being difficult to measure in an inexpensive way, the literature shows that it does not perform well in dynamic office environments [5] and does not model people's adaptability to thermal comfort [6]. We propose to measure the thermal comfort directly and thus bypass the need to estimate the comfort using indirect sensor inputs. This also allows us to simplify the user input to three values: cold (-1), neutral (0), and hot (1).

B. Comfort Data Fusion

It is a well known fact that people have different thermal comfort preferences and larger groups of people find it more difficult to agree on the optimal temperature in an office

Algorithm 1 Pseudo code of our HVAC control algorithm.

Parameters: Step, DriftStep, DriftEnabled

Input: SetpointTemp, OutdoorTemp, Votes, Occupancy

Output: HVAC_Command

function CONTROL(Run every T minutes)

if Occupancy == Empty **then**

return HVAC_PowerOff()

end if

 netVote $\leftarrow$ sum(Votes)

if netVote < 0 **then**

 SetpointTemp = SetpointTemp - Step

else if netVote > 0 **then**

 SetpointTemp = SetpointTemp + Step

else ▷ Net vote is stable

if DriftEnabled **then**

if SetpointTemp > OutdoorTemp **then**

 SetpointTemp = SetpointTemp - DriftStep

else

 SetpointTemp = SetpointTemp + DriftStep

end if

end if

end if

return HVAC_TempSet(SetpointTemp)

end function




Fig. 3. Thermal comfort of a group of people at two different set-points. Maximizing the number of comfortable people can leave many people uncomfortable (plot a), while minimizing the average discomfort finds a compromise (b).

setting. As an example, Figure 7 shows comfort zones of eight users in our building that we calculated using empirical data.

There are two basic approaches to satisfy comfort preferences of a group of people. We can maximize the number of comfortable people without regards for the comfort of the rest of the group. Alternatively, we can find a compromise between people that are uncomfortable by finding a temperature that minimizes the overall discomfort, at the possible expense of decreasing the number of people that are actually comfortable.

Figure 3a shows a hypothetical case when maximizing the number of comfortable people leads to suboptimal results. We show the temperature set-point with a horizontal line and color-code the user comfort zones by blue, green, and red, for users that feel cold, comfortable, and hot, respectively. The

control algorithm has two options. It can either satisfy the three people at lower temperatures or the three people at higher temperatures, leaving three people extremely uncomfortable. The strategy shown in Fig. 3b finds a compromise that satisfies only one person, but does not leave anyone at an extreme discomfort.

Equalizing the extremes has another advantage in our setting. Recall that our algorithm does not actually model the comfort zones of people, i.e., it would only see the number of people that are cold or hot (blue or red in Fig. 3), not how much discomfort they experience. Balancing the office temperature to achieve equilibrium can be done iteratively without quantifying user comfort preferences.

Consequently, our algorithm first determines the overall comfort of a group of users by summing their recent comfort votes, and then changes the indoor temperature by a fixed value (called *Step*) to compensate for the discomfort (see Alg. 1). This algorithm minimizes thermal discomfort of a group, without having to quantify the discomfort of each individual.

C. Temperature Step Size

One of the important parameters of our HVAC control is the value of *Step*, which defines by how much we change the temperature set-point. Clearly, a small step will cause the system react slowly to changing thermal preferences and might cause unnecessary discomfort. Similarly, if the step is too large, the system might fail to find a good compromise. For example, if *Step* is larger than 2°C in Fig. 3, the system may oscillate between two extremes and never converge to the optimal value.

Due to the limited information that we receive from users, it is difficult to estimate the optimal value of *Step* in real time. In addition, the step needs to consider system limitations, such as the maximum rate of cooling and heating that HVAC supports and the thermal inertia of the building. We treat *Step* as a system parameter that can be fine-tuned during the HVAC operation. We note that enforcing a more complex user input (e.g., on the 7-point PMV scale) would help with determining the optimal value of *Step*. However, the user input would have to be normalized to different comfort perceptions of individual users. By using the 3-point comfort scale, we sacrifice some information content in user surveys, but benefit from the input simplicity and avoid modeling of subjective user preferences.

D. Maximizing Energy Savings

Our HVAC control algorithm places people directly in control of the temperature in their air-conditioning zone based on their thermal comfort. However, our goal is also to maximize energy savings related to HVAC by adapting to changing occupancy and thermal comfort patterns in real time.

One way to save energy is to turn the HVAC off if an office is unoccupied. This approach saves energy by starting the HVAC when people arrive at work and by turning it off when people leave for lunch or meetings. Our phone

application infers user occupancy implicitly from the user data (whether people vote) and the control algorithm adapts to changing occupancy in real time. The adaptive behavior of the control algorithm also saves energy over systems with a fixed temperature setting as it can operate at a lower (or higher) temperatures based on occupant tolerances.

We achieve additional energy savings by relying on a simple yet powerful concept. We let the indoor temperature drift towards the outdoor temperature as long as it does not impact the thermal comfort of occupants. We do this by adjusting the HVAC set-point temperature; we therefore are not concerned with the actual indoor temperature.

For example in Fig. 3, the temperature can vary by 0.5°C around the optimum (24°C) without changing distribution of the uncomfortable people. Figure 3b shows the optimal temperature set-point of 24.5°C that maximizes energy savings in the summer.

We call this strategy **Drift**. The algorithm has one important parameter, *DriftStep*, a temperature by which we change the set-point when drifting towards ambient. *DriftStep* is set at a lower value than *Step* as we want to prevent the system from oscillating around the optimal temperature, which would require frequent user input. We note that one important advantage of the **Drift** algorithm is that it operates to favour energy savings both in Summer (where the goal is to minimise cooling) and in Winter (where minimal heating is desired).

One way to optimize our algorithm would be to stop **Drift** once the optimal temperature is found and restart it only if the office occupancy or thermal preferences change. We leave this for future work and show in Sec. IV that **Drift** achieves large energy savings at a relatively minor increase in the number of votes.

IV. EVALUATION

We base our evaluation on one wing of an office building at our campus. The office wing includes 20 people located in 12 offices and one conference room (Fig. 5) and there are two air-conditioning zones, one at each side of the building. It is difficult to evaluate HVAC performance of different control algorithms experimentally due to the complex logistics of such experiments and day-to-day changes in office occupancy. Instead, we built a simulator that can replay empirical data from historical traces, including office occupancy, indoor/outdoor temperatures, thermal comfort of people, voting through a smart-phone application, and thermal response of building to HVAC actuation. We first describe the simulator, then evaluate performance of our control algorithm, and finally compare its performance to existing fixed-set-point HVAC control strategies.

A. Simulation Engine

We built a simulator that uses empirical data to simulate different HVAC control strategies in realistic scenarios. To provide data for the simulator, we collected empirical data from one of our buildings. We measured occupancy patterns and thermal preferences of people, values of indoor and




Fig. 4. Climate node for temperature and occupancy sensing. Fig. 5. Floorplan and sensor locations in our building.

outdoor environmental parameters, thermal response of the building, and thermal characteristics of the HVAC system. This section provides a brief overview of the simulator and more details can be found in [15].

1) *Empirical Data*: We deployed a wireless sensor network to measure environmental parameters in our building. Data from this network were used as the ground truth for our simulations. We call our mote platform a *personal climate node* (see Fig. 4). The platform was developed in-house around a low-power Intel 8051 MCU and sub 1-GHz Nordic radio. The climate node includes passive infrared (PIR), temperature, and humidity sensors. The sensors self-organize in a wireless mesh network and transmit data to a base station with one minute period. We use office temperature measurements to estimate thermal comfort of people and PIR sensors to estimate office occupancy. We placed climate nodes at each person's desk to minimize the rate of false positives of PIR sensors. We collected the data for a period of one year, capturing environmental parameters during different weather seasons and work schedules.

We also installed a PC application for surveying people about their current thermal comfort on the PCs of all participants. The application allows us to trigger surveys remotely and saves responses in a database. We administered the surveys over a three week period, asking participants to rank their current comfort on the ASHRAE 7-point scale. Simultaneously, we were changing the HVAC temperatures across a wide range, including uncomfortably hot and cold, so that we could build detailed thermal comfort models for each individual.

Finally, we obtained access to our building control and management system for our HVAC. We were logging temperature set points and duty cycles of heaters and air conditioning units for all HVAC zones under study. Additionally, we changed the indoor temperature through oBIX [1], a RESTful Web-based interface to building control systems. We used the HVAC control to obtain comprehensive data on users' comfort levels in the survey trials, but also to characterize the thermal response of our building to different HVAC settings.

2) *Empirical Models*: For each participant, we built a thermal comfort model by cross-referencing the office temperature data from climate domes and thermal comfort surveys. A response of -3 indicates a high level of discomfort due to




Fig. 6. Example of two users with different thermal preferences.




Fig. 7. Green lines show comfort zones of eight users in our building. We inferred the zones from empirical data.

cold, while a value of +3 indicates a high level of discomfort due to heat. Interviews with participants demonstrated that a response in the range $[-1, +1]$ was indicative of comfort, see an example for two participants in Fig. 6. Data points from survey responses are shown in blue. After some basic filtering for outliers, we estimate the comfort limits of each participant as red crosses in the figure. We show the resulting comfort models for all our users with valid comfort data in Figure 7. We do not consider the remaining users for whom there was insufficient or inconsistent survey data. The models show a high degree of correlation between the preferences of individual participants. This observation suggests that it is possible to deliver an internal temperature that is comfortable for the majority of participants. Note that this modeling is required for evaluation purposes only, for example, to simulate thermal comfort of people at a given temperature. The HVAC control algorithm presented in this paper remains model-free, and the sensing apparatus described above does not form part of a final deployment.




Fig. 8. Different thermal response of a building on different days.




Fig. 9. Indoor temperature for different temperature step size.

Our simulator also models thermal response of our building at different HVAC settings. We use empirical data to model the rate of indoor temperature change given the outdoor temperature and the value of HVAC set-point, the duty-cycle of HVAC given the indoor-outdoor temperature difference, and decay of indoor temperature towards the ambient when HVAC is turned off (see [15] for more details). These models allow us to estimate the indoor temperature in our building and the energy usage of HVAC under different control strategies. We note that such analysis eliminates the need for any sophisticated modelling of building materials or infrastructure, yet is capable of capturing complex temporal behaviour of the building (see Figure 8) and provides a basis for realistic evaluation of different building control algorithms. The fixed control algorithm used as a baseline is evaluated within this framework.

3) *Simulation*: The simulator uses comfort models derived for each participant and generates smart-phone user comfort input given the current temperature in the building. We simply replay the historical data to simulate occupancy of individual offices. The indoor temperature is simulated based on the outdoor ambient temperature and the actions of the simulated HVAC control algorithm. We evaluate 5 algorithms: fixed temperature set-point algorithm at three levels (*Fixed 21.5, 23, and 24.5*), our adaptive model-free control algorithm (*Ambient*), and our adaptive algorithm with ambient drift enabled (*Drift*). The current control algorithm used in our building is *Fixed 23*.

B. Model-free Data Fusion

In this section, we discuss the selection of system parameters and the impact of their misconfiguration on HVAC performance. We also evaluate the performance of our model-free HVAC control algorithm in terms of adaptivity to different weather seasons and different occupancy patterns.

1) *Step size*: We study the optimal setting of the *Step* system parameter in a series of experiments in Fig. 9. Clearly, if the step size is too small, the system is too slow in adapting to changed conditions and the indoor temperature takes a long time to reach comfortable temperatures. If the step is too large, however, the system is not be able to find a comfortable




Fig. 10. Adaptive HVAC control with drift towards outdoor ambient temperature. Ambient drift eventually causes the temperature to fall outside the comfort zone of a person, which results in temperature being reset back by the control algorithm.

temperature and oscillates around the temperatures that would make majority of occupants comfortable. Another disadvantage of the large step is that it does not let the control algorithm to adjust the indoor temperature within the comfortable zone. In the absence of the ambient drift technique, this might result in high energy usage of HVAC.

We found that the value of 1°C worked well in practice in our deployments. The value is small enough to find temperatures that make participants comfortable, while it allows a fast response when office occupancy changes. Interestingly, the rate of change of 1°C per 30min is similar to a typical rate of outdoor ambient temperature change in our climate zone which suggests that this rate of change is naturally accepted by people.

2) *Adaptiveness*: We ran the *Drift* version of our algorithm on two days in different seasons and show that it adapts well to different ambient temperatures (see Figure 10). We can clearly see the two steps working together to find the optimal temperature. First, the larger *Step* of 1°C is used to reach comfort zones of people fast and then the *DriftStep* towards the ambient temperature to find energy-optimal point within the comfort zone. As we have already noted, a promising avenue for future work would be to detect and pre-empt this step in order to establish a stable temperature in the absence of occupancy changes.

Eventually, the ambient drift will push the indoor temperature outside the comfort zones of people. This results in people providing additional feedback and resetting the temperature back in the comfort zone. Careful analysis of plots in Fig. 10 shows that the temperature set-points used during the colder day are lower on average than those on the warmer day, which demonstrates that our HVAC control algorithm automatically adjusts to different weather seasons.

To evaluate adaptiveness of our algorithm to changing occupancy patterns we turned off the *Drift* mechanism and show the indoor temperature in Fig. 11. We also plot the office occupancy on the secondary y axis. We see that system reacts at 7am by increasing indoor temperature in the morning after the first person arrives at work. The system turns off the HVAC while the person prepares morning tea at 8am as shown by the indoor temperature drifting towards the ambient. The temperature reaches a comfortable level after 9am, when the majority of office occupants arrive at work. The system




Fig. 11. The vote-based HVAC control adapts to changing occupancy.

turns HVAC off second time during the lunch break, which again results in higher indoor temperature due to the drift towards ambient. As most occupants are comfortable at this temperature, the system does not have to change it until people leave work at 6pm. This second stable set-point is closer to the ambient temperature, and thus saves more energy than would be possible without considering occupancy.

C. Comparison to Fixed Temperature Control

This section explores the impact of our approach on HVAC energy consumption and occupant comfort. We evaluate both Adaptive and Drift variants of our approach and compare the results to control algorithms that maintain fixed temperature. We use empirical data collected during 3 weeks in the spring, where the temperature in the morning is generally below the comfortable zone of most people, but it frequently goes above that comfort zone during the day. Thus the HVAC control needs to both heat (morning and evening) and cool (day) the building throughout our experiments.

We evaluate energy consumption of HVAC in kWh and also in terms of $^{\circ}\text{C}$ as a mean difference between the indoor and ambient temperatures. The intuition is that the larger the difference of indoor and outdoor temperature, the higher the energy consumption. In practice, however, air-conditioning and heating use different amount of energy and thus the relationship might not hold in all weather seasons. We also evaluate user comfort in terms of the deviation of the indoor temperature from their comfort zone in $^{\circ}\text{C}$ and count the number of votes that users provide on average.

Figure 12 compares the energy consumption of our approach with fixing HVAC set points at a given level throughout the day. Both flavors of our approach reduce energy consumption over fixed set points in terms of $^{\circ}\text{C}$. Due to heating being cheaper than air-conditioning in our building, the *Fixed 24.5* strategy actually uses less energy in terms of kWh than the *Adaptive* algorithm. However, the *Drift* approach saves a significant amount of energy compared to all fixed strategies in terms of both kWh and $^{\circ}\text{C}$. Nearly 50% and 65% kWh of energy can be saved over the most energy efficient (24.5) and




Fig. 12. Average daily energy consumption: voting-controlled HVAC reduces energy usage by maintaining the indoor temperature closer to the ambient temperature. The horizontal split on the energy consumption plot indicates energy used due to air-conditioning (top) and heating (bottom).




Fig. 13. Average daily discomfort over 3 weeks: both variants of voting-based HVAC control deliver comfort performance that is only a fraction of a degree different from optimal comfort levels.

Fig. 14. Voting requirements: drift approach requires slightly more votes.

the most comfortable (23) fixed techniques.

Drift achieves its large energy reductions by exploiting situations when the indoor temperature hovers around the comfort equilibrium of occupants, and then switching off the HVAC while the temperature drifts toward ambient. Once that drift starts affecting comfort equilibrium, the HVAC is then turned on again. This approach is particularly attractive to commercial HVAC systems where switching cooling or heating from a particular zone is an ON/OFF process and does not involve significant mode switching costs because the central system remains active throughout the day.

Figure 13 shows the resulting average deviation from comfort zones for all occupants in our building. While the fixed set point approach with 21°C leads to about 0.7 degree deviation from comfort zones on average, all other approaches appear to cause between 0.15 and 0.3 $^{\circ}\text{C}$ deviation. We suggest that this is practically imperceptible. Thus, our approach delivers significant HVAC energy savings with nearly no impact of occupant comfort.

A remaining question about our approach is the extent of involvement that it requires from participants. Figure 14 shows that the average number of votes for our approach ranges between 4 and 6 votes per occupant per day, which we believe is minimally inconvenient to users of the system. In particular, the drift strategy requires approximately 1 vote per occupant more per day than the adaptive strategy which is well worth

the achieved energy savings.

D. Evaluation Summary

Based on empirical data gathered from one of the buildings on our campus, we were able to build thermal comfort models of occupants and estimate various environmental parameters. Using this information, we developed a simulator capable of comparing the effect (inter alia, in terms of energy savings and occupant comfort) of different control algorithms when seeded with real data. The simulator was constructed such that algorithms tested in it could be used without alteration as control algorithms in the live system.

We focused our analysis on comparing three fixed set-point control schemes – *Fixed 21.5*, *Fixed 23*, and *Fixed 24.5* – to two adaptive algorithms – *Adaptive* and *Drift*. Results from our simulator showed that both flavours of our algorithm reduced energy consumption relative to fixed set-points in terms of °C. In the case of our building, the *Drift* approach saved 50-65% energy compared to the fixed algorithms, with a negligible impact on comfort.

This demonstrates the key advantage of our approach: by utilising participatory sensing, we are able to obtain significant energy savings within the bounds of occupant comfort. Moreover, we can achieve this result without the effort or expense of modelling comfort directly.

V. DISCUSSION AND RELATED WORK

A. Model and Sensor Free HVAC Control

Our results show that model and sensor free HVAC control can achieve at least 50% reduction in energy consumption with minimal impact on thermal comfort. The participatory sensing aspect of our approach is inspired by Thermovote [8], which also uses human input as the primary information stream for control decisions. Our system, however, uses only human input for HVAC control, while Thermovote estimates average P.M.V parameters for all occupants (except temperature, which is measured), which is prone to errors resulting from individual variations all these parameters.

In addition, our system requires only knowledge of the actual (or expected) outdoor temperature in order to determine a set-point. We do not measure the actual indoor temperature or determine an offset based on actual mean vote, but rather decide whether to increase or decrease the current set-point on each iteration. This removes an important source of complexity and potential error.

Finally, users of our system only need to indicate if they are hot or cold rather than voting on the 7 point ASHRAE scale, as is the case with Thermovote. A desirable feature in Thermovote is its real-time calculation of the offset temperature, which in our approach is the equivalent of adaptively computing the step change in HVAC temperature set point every cycle. We leave the exploration of this feature for future work.

B. Voting Resolution

While moving away from the 7 point thermal comfort scale provides simplicity, it also has inherent limitations in its coarse-grained information content. It is not possible to distinguish occupants that are slightly uncomfortable from other occupants that may be extremely uncomfortable. The current implementation only considers the most recent vote of a user in any voting period. One option to extract information on the severity of thermal discomfort is to consider the frequency of votes, with the assumption that occupants who are extremely uncomfortable will tend to vote more often. This additional information would need to be balanced against the obvious potential for abuse. Such a feature could also act as a tiebreaker in deadlock voting scenarios where there is an equal number of occupants who vote cold and hot. Currently, our algorithm either maintains the current set point or drifts towards ambient when the number of hot and cold occupants is equal, as it aims to minimise the discomfort among all occupants. Fine-tuning this aspect of the algorithm offers a clear opportunity for improvement.

C. Learning Models

A major advantage of our approach is the absence of any individual level parameter estimation, avoiding errors that can arise from differences across people. While the use of individual level models based on empirical data [4], [14] has been shown to improve HVAC control, building these individual level models is typically a labor-intensive and potentially intrusive process for participants. An interesting area for future work is to learn occupant comfort models over time, by inferring the limits of their comfortable temperature zones from their historical thermal comfort input. This would bring models back to our approach based on empirical data rather than blanket estimations of individual level parameters. One important advantage of this approach would be to minimise the number of user votes required, while still retaining the advantages of on-demand participatory sensing.

D. Privacy, Security and Incentives

Privacy and incentives are key considerations for any participatory sensing approach. Erickson et al. [8] indicate that participants often readily contribute their thermal comfort information on the 7 point ASHRAE scale, since this information is not considered sensitive. Our HVAC control approach requires only relative thermal comfort input from participants (whether they are cold or hot), and avoids information capture on their personal attributes, such as metabolism rate or clothing insulation [14].

In terms of incentives, participants in our approach are genuine stakeholders as the outcome of the HVAC voting process will affect their thermal comfort. As such, we expect that building occupants will have an intrinsic interest in voting proactively when they experience thermal discomfort.

The use of participatory sensing to control a physical system involves potential security issues, with the potential for malicious users to attempt to steer the system toward an

undesirable state. For instance, malicious users may try to keep voting that they are too hot on a summer day to keep pushing the HVAC setpoint lower, which can cause both occupant discomfort and higher energy consumption. We believe the majority of malicious voting patterns can be detected and filtered out at the central controller to avoid such situations arising.

VI. ACKNOWLEDGEMENTS

We would like to acknowledge the help of Luke Hovington and Teddy Shih in helping us to maintain the wireless sensor network to collect data from climate domes and power meters. We also value patience of the CSIRO ICT Centre's Pervasive Computing group during short periods of extreme temperatures in our HVAC trials and thank them for spending the time to submit thermal comfort surveys.

REFERENCES

- [1] C. Aeron, *Obix*. Onym Press, 2013. [Online]. Available: <http://books.google.com.au/books?id=5ArdlwEACAAJ>
- [2] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng, "Duty-cycling buildings aggressively: The next frontier in hvac control," in *IPSN*, X. D. Koutsoukos, K. Langendoen, G. J. Pottie, and V. Raghunathan, Eds. IEEE, 2011, pp. 246–257.
- [3] ASHRAE, "ASHRAE standard 55: Thermal environmental conditions for human occupancy," in *ASHRAE, Inc*, 2010.
- [4] R. de Dear, "A global database of thermal comfort field experiments," in *ASHRAE Trans*, vol. 104, no. 1, 1998.
- [5] R. de Dear and A. Auliciems, "Validation of the predicted mean vote model of thermal comfort in six Australian field studies," in *ASHRAE Transactions*, vol. 91, no. 5, pp. 452–468, 1985.
- [6] R. de Dear and G. Brager, "The adaptive model of thermal comfort and energy conservation in the built environment," *International Journal of Biometeorology*, vol. 45, no. 2, pp. 100–108, 2001.
- [7] V. L. Erickson, M. Carreira-Perpiñán, and A. E. Cerpa, "Observe: Occupancy-based system for efficient reduction of hvac energy," in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*. IEEE, 2011, pp. 258–269.
- [8] V. L. Erickson and A. E. Cerpa, "Thermovote: participatory sensing for efficient building hvac conditioning," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ser. BuildSys '12. New York, NY, USA: ACM, 2012, pp. 9–16. [Online]. Available: <http://doi.acm.org/10.1145/2422531.2422534>
- [9] L. J. and T. A., "Sensing data centres for energy efficiency," *Philos Transact A Math Phys Eng Sci*, vol. 370, no. 1958, pp. 136–157, 01 2012. [Online]. Available: <http://pubget.com/paper/22124086>
- [10] X. Jiang, M. Van Ly, J. Taneja, P. Dutta, and D. Culler, "Experiences with a high-fidelity wireless building energy auditing network," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New York, NY, USA: ACM, 2009, pp. 113–126. [Online]. Available: <http://doi.acm.org/10.1145/1644038.1644050>
- [11] A. Kamthe, V. Erickson, M. A. Carreira-Perpiñán, and A. Cerpa, "Enabling building energy auditing using adapted occupancy models," in *Proceedings of the Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ser. BuildSys '11. New York, NY, USA: ACM, 2011, pp. 31–36. [Online]. Available: <http://doi.acm.org/10.1145/2434020.2434028>
- [12] R. H. Katz, D. E. Culler, S. Sanders, S. Alspaugh, Y. Chen, S. Dawson-Haggerty, P. Dutta, M. He, X. Jiang, L. Keys, A. Krioukov, K. Lutz, J. Ortiz, P. Mohan, E. Reutzel, J. Taneja, J. Hsu, and S. Shankar, "An information-centric energy infrastructure: The berkeley view," *Sustainable Computing: Informatics and Systems*, vol. 1, no. 1, pp. 7 – 22, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2210537910000028>
- [13] 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 *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 211–224. [Online]. Available: <http://doi.acm.org/10.1145/1869983.1870005>
- [14] B. W. Olesen and K. C. Parsons, "Introduction to thermal comfort standards and to the proposed new version of en iso 7730," in *Energy and Buildings*, 2002.
- [15] S. Purdon, B. Kusy, and R. Jurdak, "Empirically driven simulation of people in urban environments," CSIRO, ICT Centre, CSIRO, Australia, Tech. Rep. CSIRO-xxx, May 2013.
- [16] A. Schoofs, D. T. Delaney, G. M. P. O'Hare, and A. G. Ruzzelli, "Copolan: non-invasive occupancy profiling for preliminary assessment of hvac fixed timing strategies," in *Proceedings of the Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ser. BuildSys '11. New York, NY, USA: ACM, 2011, pp. 25–30. [Online]. Available: <http://doi.acm.org/10.1145/2434020.2434027>
- [17] J. Wall, J. K. Ward, and S. West, "Autonomous controllers for intelligent hvac management," in *World Sustainable Building Conference*, 2008, pp. 1259–1267.