

# SPOT: A Smart Personalized Office Thermal Control System

Peter Xiang Gao and S. Keshav  
School of Computer Science  
University of Waterloo, Waterloo, Ontario, Canada

## ABSTRACT

Heating, Ventilation, and Air Conditioning (HVAC) accounts for about half of the energy consumption in buildings. HVAC energy consumption can be reduced by changing the indoor air temperature setpoint, but changing the setpoint too aggressively can overly reduce user comfort. We have therefore designed and implemented SPOT: a Smart Personalized Office Thermal control system that balances energy conservation with personal thermal comfort in an office environment. SPOT relies on a new model for personal thermal comfort that we call the Predicted Personal Vote (PPV) model. This model quantitatively predicts human comfort based on a set of underlying measurable environmental and personal parameters. SPOT uses a set of sensors, including a Microsoft Kinect, to measure the parameters underlying the PPV model, then controls heating and cooling elements to dynamically adjust indoor temperature to maintain comfort. Based on a deployment of SPOT in a real office environment, we find that SPOT can accurately maintain personal comfort despite environmental fluctuations and allows a worker to balance personal comfort with energy use.

## Categories and Subject Descriptors

H.4.0 [Information Systems Applications]: General

## General Terms

Design, Human Factors

## Keywords

HVAC, Personalization, Human Thermal Comfort

## 1. INTRODUCTION

About 30% to 50% of the residential and commercial energy consumption in most developed countries is used by Heating, Ventilation, and Air Conditioning (HVAC) systems [3, 5, 22, 25]. Increasing the efficiency of HVAC systems, therefore, can greatly reduce the overall energy footprint of a commercial building.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*e-Energy'13*, May 21–24, 2013, Berkeley, California, USA.  
Copyright 2013 ACM 978-1-4503-2052-8/13/05 ...\$15.00.

The focus of our work is thermal comfort in office environments. We assume that workers in offices have work areas that are relatively thermally isolated from each other, such as separate offices or cubicles with walls. Thus, heating and cooling within a personal work space would be for the benefit of a single worker.

We suggest that the overall building temperature level be set to a value lower than normal in winter and to a value higher than normal in summer. Then, a personal thermal controller in each work space could provide an offset to this base temperature. For instance, most commercial buildings today are heated to 23°C in winter. Instead, we suggest that the buildings be heated only to, say, 20°C, and that each work space have a small computer-controlled radiant heater that can heat the work space to a personalized higher level. In summer, symmetrically, a small fan can provide additional cooling below a building setpoint of, say, 26°C [28]. The role of the personal thermal control system, therefore, is to automatically control the per-workspace radiant heater or fan to maintain the comfort level of individual workers when they are actually present. In contrast, existing time-based or motion-based sensor control often suffers from irksome false positives and false negatives. Manual control, of course, would have no such errors, but this requires human effort, and office workers have no incentive to participate.

It has been found that user comfort is not just a function of room temperature. Two persons who are differently dressed would experience different levels of comfort for the same room temperature. Ideally, an HVAC control system should control room temperature not to achieve a temperature setpoint, but a particular human comfort level. This is the key idea that motivates the design of SPOT: a Smart Personalized Office Thermal control system.

SPOT uses an ensemble of sensors to measure the six parameters that have been found to contribute to human comfort: air temperature, radiant temperature, humidity, air speed, clothing level, and activity level. This lets it compute human comfort according to the ISO 7730 standard called the Predicted Mean Vote (PMV) model [4]. We have extended this model to allow per-user personalization; we call our personalized model the Predicted Personal Vote (PPV) model. SPOT uses the PPV model to maintain a desired comfort level despite environmental fluctuations. We have deployed SPOT and evaluated its performance in a realistic office environment. Our work makes it possible to trade off a decrease in human comfort for a reduction in energy usage.

The major contributions of our paper are:

- We extend the ISO 7730 standard [4] to define the PPV model for user comfort and use it to design SPOT, an HVAC control system that maintains user comfort, rather than merely air temperature
- We have implemented SPOT and deployed it in a realistic environment

- We find that SPOT can accurately maintain personal comfort despite environmental fluctuations and allows a user to balance personal comfort with energy use.

## 2. BACKGROUND

HVAC control systems traditionally put user comfort first, expending energy freely to achieve a given setpoint. ‘Dumb’ thermostats use the same setpoint all day, and smarter, programmable thermostats allow users to vary setpoints by time of day and day of week. Some thermostats allow remote control. For example, in Ontario, the PeakSaver [1] thermostat responds to an emergency broadcast radio signal and increases the cooling set point by up to two degrees, thereby reducing home electricity usage by up to 37%. Other ‘smart’ thermostats, such as the Nest [6], learn user occupancy patterns to intelligently control HVAC usage by means of proprietary algorithms. Nevertheless, none of these thermostats are aware of user comfort: they focus, instead, only on controlling room temperature.

The basis of our work is a quantitative model for human comfort called the PMV model that is defined in the ISO 7730 Standard [4]. The PMV model computes a numerical comfort level, called a *vote*, that describes the degree of comfort of a typical person in a moderate thermal environment. The PMV model predicts human comfort as a function of four environmental variables (air temperature, radiant temperature, air speed, and humidity) and two personal variables (clothing and physical activity). Given these variables, it predicts the mean value of a group of people’s votes in a 7-point ASHRAE [2] thermal sensation scale.

| Vote | Comfort Level |
|------|---------------|
| +3   | Hot           |
| +2   | Warm          |
| +1   | Slightly Warm |
| 0    | Neutral       |
| -1   | Slightly Cool |
| -2   | Cool          |
| -3   | Cold          |

Table 1: 7-point ASHRAE scale in PMV model

The PMV model was first proposed by Fanger [15] in 1970 and it is widely used for evaluating thermal comfort [7] [20]. Although the model is based on a theoretically well-grounded physical thermal balance model, it has been found to be problematic to use in practice [18]. Many variations of the PMV model have been developed to fix these problems. For example, De Dear et. al. [10] developed a model to capture the sociological and geographical factors that may affect human’s thermal preference, such as people living in warmer areas preferring warmer indoor temperature than people living in cooler areas. Similarly, Nicol et. al. [24] have shown that people can use physiological and psychological adaptations to be comfortable in a wider range of temperatures than supposed by the PMV model; their model reflects this observation.

Although these newer models improve the accuracy of the PMV model, they all predict the average thermal comfort of a large group of people. However, in a micro-climate such as an office work area, comfort is usually relevant only for one person or a small number of people. This motivates the design of a *personalized* thermal comfort model. In our work, we extend the PMV model to the Predicted Personal Vote (PPV) model to capture individual thermal preference. We use PPV model to automatically adjust an HVAC

control system’s temperature setpoint so that a worker always feel comfortable.

### 2.1 Predicted Mean Vote

We now describe the PMV model [15] in greater detail. It assigns a numerical comfort value  $pmv(\mathbf{x})$  based on a vector  $\mathbf{x}$  with six elements

$$\mathbf{x} = \{t_a, \bar{t}_r, v_{ar}, p_a, M, I_{cl}\}^T$$

- $t_a$  is the air temperature
- $\bar{t}_r$  is the mean background radiant temperature
- $v_{ar}$  is the air velocity
- $p_a$  is the humidity level
- $M$  is the metabolic rate of a person
- $I_{cl}$  is the clothing insulation factor of a person

We can evaluate PMV using the function:

$$pmv = pmv(\mathbf{x}) \quad (1)$$

The details of the function can be evaluated in practise are in the appendix.

## 3. DESIGN

We now describe our design in more detail. Recall that SPOT’s goal is to maintain a particular comfort level (PPV value) based on sensor measurements and its control over the operation of a small personal radiant heater or fan. We first discuss the PPV model and clothing level estimation, then SPOT’s control strategy.

### 3.1 Predicted Personal Vote Model

The PMV model reflects the thermal comfort of a large group of people. However, individual workers may have their own thermal preference. We have, therefore, modified the PMV model to create a model we call the Predicted Personal Vote (PPV) model.

For each person, the Predicted Personal Vote function has two parts, the PMV part and the personal part:

$$ppv(\mathbf{x}) = pmv(\mathbf{x}) + personal(\mathbf{x}) \quad (2)$$

where  $pmv(\mathbf{x})$  is the output of the PMV model and  $personal(\mathbf{x})$  models how the current user is different from an average person. We model  $personal(\mathbf{x})$  as a linear function:

$$personal(\mathbf{x}) = \mathbf{a}^T \mathbf{x} + b \quad (3)$$

where  $\mathbf{a}$  is a vector of size 6 that models the users sensitivity to each variable.

$$\mathbf{a} = \{a_{temp}, a_{radiant}, a_{velocity}, a_{humidity}, a_{metabolic}, a_{clothing}\}^T \quad (4)$$

. For example, a person who is more sensitive to humidity than average will have a relatively large  $a_{humidity}$  value. Variable  $b$  denotes the thermal preference of the user. A person prefers warmer temperatures will have negative  $b$  value and vice versa.

Using the PPV model requires a training phase. In the training phase, SPOT measures the environmental variables  $\mathbf{x}$  and also records the worker’s personal vote  $apv$ . Suppose we have a training set  $\{(\mathbf{x}_k, apv_k)\}_{k=1}^K$  of size  $K$ , where  $apv_k$  is the  $k$ -th actual personal vote, and  $\mathbf{x}_k$  is the vector of environmental and personal variables when the user gives the  $k$ -th vote. This allows us to estimate parameters  $\mathbf{a}$  and  $b$  using straightforward linear regression. In the absence of a training set, SPOT simply reverts to the PMV

model. Similarly, when there are not enough data points to do a linear regression for Equation 3, we train a simpler linear function  $g(\cdot)$  to estimate PPV:

$$ppv(\mathbf{x}) = g(pmv(\mathbf{x})) \quad (5)$$

The function  $g(\cdot)$  is trained by least square regression.

### 3.2 Clothing Level Estimation

Five out of the six underlying parameters of the PPV model can be measured in a relatively straightforward manner using appropriate sensors (this is discussed in more detail in §4.2). However, measuring the ‘clothing level’ parameter is non-trivial (see Table 12), and the focus of this subsection.

The key idea behind our approach to clothing level estimation is the fact that most humans have a relatively constant skin temperature of about 34°C. The greater the level of clothing worn, the greater the degree of insulation, and the lower the temperature of the outermost layer of clothes. Thus, the clothing level can be estimated by measuring the temperature of the clothing using an infrared sensor as discussed in §4.2.2.

Specifically, we build a linear regression model to estimate the clothing level  $I_{cl}$  as:

$$I_{cl} = f(t_{ir}) \quad (6)$$

where  $f(\cdot)$  is a linear function and  $t_{ir}$  is the infrared intensity of the clothing. We fit the function  $f(\cdot)$  using least square linear regression. The model is trained using a data set of  $I_{cl}$  estimates from Table 12 and  $t_{ir}$  measured by the infrared sensor.

Note that this assumes that the worker’s body temperature is in the normal range. In case the worker has a fever, estimation accuracy can be affected. We can solve this problem by using the infrared intensity of the worker’s face as a reference; however, we have not currently implemented this refinement.

### 3.3 Control Strategy

SPOT maintains the PPV of a worker by controlling the air temperature, because this is the factor underlying the PPV that is the easiest to control. It uses a simple reactive control strategy rather than a complex model-based predictive approach, such as the one in [7]. This is possible because personal heating and cooling systems such as radiant heaters and fans affect human comfort almost immediately<sup>1</sup>, unlike centralized HVAC systems, such as air conditioners and forced-air heaters, which can take tens of minutes to take effect.

Specifically, when the Kinect sensor indicates the presence of a worker in the work space during the prior five minutes, the system chooses an operative temperature setpoint such that  $ppv(\mathbf{x}) = dc$ , where  $dc$  is the desired comfort level (nominally 0). The five-minute window allows thermal comfort to be maintained despite brief absences. For example, in winter, when a worker is detected and  $ppv(\mathbf{x}) < dc$ , the heater is turned on to increase the room temperature. Otherwise, the heater is turned off.

This reactive control strategy takes into account both real-time occupancy and personal thermal comfort. An occupancy-aware reactive controller will always use less energy than an occupancy-unaware controller. On the other hand, a worker may choose a personal comfort value that is much higher than normal, causing the corresponding heater or fan to expend more energy than normal. However, we believe that this would be more than made up by the reduction of energy use in common areas that are heated or cooled to lower or higher than a nominal setpoint, respectively.

<sup>1</sup>We validate this observation in §5.3.




Figure 2: The image at left is from the camera and the image at right is the depth image. Lighter portions of the image are closer to the camera; the green portion of the image is too far to be measured, and there is no depth information available from black portions of the image.

## 4. IMPLEMENTATION

We now describe the implementation of our system in greater detail. SPOT has three principal components: controller, sensors and actuators.

### 4.1 Controller

The SPOT controller is a PC with an Intel i5-3450 processor and 8GB of RAM, running Windows 7 Enterprise edition. The entire project code is written in C#. All sensors and actuators are connected to this PC, and all control logic is implemented on this machine.

### 4.2 Sensors

SPOT uses multiple sensors to measure the environmental variables (air temperature, background radiant temperature, humidity and wind speed) and personal variables (users’ clothing level and activity level) that underlie the PPV model. We describe these next.

#### 4.2.1 Microsoft Kinect

A Microsoft Kinect sensor provides 3D information in real-time about the location of humans in a scene. The Kinect sensor was originally designed for the Microsoft Xbox as a natural user interface. By using a Kinect sensor, video game players can interact with Xboxes without actually touching the game controllers.

The Kinect has an RGB camera and an infrared camera. The RGB camera is similar to a normal webcam that captures image from the visible light spectrum. The infrared camera works together with an infrared projector, which emits infrared laser signal with a predefined pattern. The infrared camera collects the reflected laser beams and calculate the distance to the laser point by the time difference between sending and receiving the signal. The Kinect can generate  $640 \times 480$  resolution depth images with a sensitivity of 1mm. Figure 2 shows the raw infrared image and the depth image generated by the Kinect sensor.

By using both the color images generated by the RGB camera and the depth images generated by the infrared camera, the Kinect can build a 3D motion model for the player. When a player enters the frame, the Kinect starts to track the skeleton points of the player, and report the locations of these skeleton points (such as head, hands, knees) as a *skeleton stream*. Gesture-based Xbox applications can use the skeleton stream as user input.

For our research prototype, we use Kinect for Windows, which is a special sensor designed for Windows developers. We implemented our system using Visual Studio 2010 and Kinect for Windows SDK v1.6.

SPOT uses the Kinect for three purposes:




Figure 3: Infrared Thermometer and WeatherDuck Climate Monitor

1. It is used as an *occupancy sensor*. When a person is tracked in the skeleton frame, the work space is treated as occupied and the system starts to control the room temperature. We also use the Kinect skeleton tracking APIs to determine the worker's activity level.
2. It is used as a worker *location sensor*. We use the location information to point an infrared thermal sensor mounted on a tracking system at the worker to measure the worker's clothing surface temperature. This allows us to estimate the clothing level (see §3.2).
3. We also use the Kinect to allow workers to *customize their PPV model* parameters using simple gestures. To record a vote, a worker simply points to the Kinect and raises his or her hand in the air to indicate a particular comfort level (the selected comfort level is shown on a screen connected to the Kinect). The system then records one data point of the form  $(x_k, apv_k)$  (see §3.1).

#### 4.2.2 Infrared Thermometer

A MLX90614 Infrared Thermometer (Figure 3 upper part) detects background radiant temperature between  $-40^{\circ}\text{C}$  to  $+85^{\circ}\text{C}$  with a resolution of  $0.02^{\circ}\text{C}$ . It is connected to an Arduino Uno<sup>2</sup> board, which reads the measured radiant temperature and sends the value to a PC via a USB cable every second.

To measure the surface clothing temperature, we mounted two servos<sup>3</sup> and an infrared sensor on top of the Kinect (Figure 4). The infrared sensor and a laser pointer are placed on the two servos such that they can face any direction. The laser pointer is used for calibration, and turned off during normal operation. The two servos, the infrared sensor, and the laser pointer are connected to an Arduino micro-controller. The micro-controller sends signals to control the angle of the servos and the on/off state of the laser pointer. The micro-controller is connected to the PC with a USB cable, and it communicates with the PC program using a virtual serial port.

When a worker enters the work space, the Kinect tracks the worker and sends a skeleton stream to the PC. The PC finds the location of the worker's body center and calculates the rotation angle

<sup>2</sup><http://arduino.cc>

<sup>3</sup>A servo is similar to a stepper motor in that its degree of rotation can be precisely controlled.




Figure 4: Kinect with infrared sensor mounted. The infrared thermal sensor and laser pointer are installed on top of the two servos, which can adjust the rotation angles of the infrared sensor and laser pointer. The micro-controller controls the laser pointer and servos. It also pulls infrared readings from the sensor.

of the servos. It then communicates with the micro-controller to adjust the angle of the two servos so that the infrared sensor is facing the body center. When the tracked worker is moving, the infrared sensor may not be actually facing towards the worker. Therefore, we introduce a 0.5 second measurement delay into the system. That is, the infrared sensor starts collecting data only when the worker has been standing still for at least 0.5s. The system then estimates the clothing insulation by the clothing surface temperature as described in §3.2.

#### 4.2.3 Environment Sensor

SPOT senses environmental variables using the WeatherDuck Climate Monitor<sup>4</sup> (Figure 3 lower part), a low-cost sensor that monitors air temperature, humidity and air flow. It can detect air temperature from  $-10^{\circ}\text{C}$  to  $85^{\circ}\text{C}$  and relative humidity from 0% - 100%. Its air flow sensor can detect wind speed from 0 to 100 CFM. It also measures the light and sound level of the room as side channels for occupancy detection. The WeatherDuck Climate Monitor is connected to the PC via a serial to USB converter.

### 4.3 Actuator

SPOT controls a SunBeam SLP3300CN heater with a maximum power rating of approximately 1350W using a power plug that is controlled over a Z-Wave wireless network. Z-Wave is specially designed for reliable, low-latency communication of small data packets, which is desirable for home appliance control. Z-Wave devices use command classes to achieve different tasks. In our research prototype, we use a DSC06106 Smart Energy Switch to control and sense the energy consumption of the heater. The Smart Energy Switch is controlled wirelessly by a Silicon Labs CP201s Z-Wave controller, which supports command classes to control the on/off state of a device and measure the energy consumption of that device.

### 4.4 Occupancy Detection

SPOT detects room occupancy using the Microsoft Kinect. Recall that the Kinect APIs allows the controller to obtain near-real-time skeleton tracking. When there is a skeleton tracked by Kinect,

<sup>4</sup><http://www.itwatchdogs.com>




Figure 5: Heater control logic showing the leaky-bucket based low-pass filter. An office room will be declared as occupied only if it is occupied for a certain fraction of the past few minutes.

SPOT considers the room as occupied. However, it does not turn the heater on immediately after it detects a worker to deal with transient occupancy of the work space. Instead, we have implemented a leaky-bucket based low-pass filter that turns on the heater only if the work space has been occupied for a sufficiently long fraction of the prior few minutes.

Specifically, the system has a virtual leaky bucket of size 5 units. The bucket is initially empty. At the end of each minute, if the Kinect sensor reports that the work space was occupied in the past minute, one unit of “water” is added to the leaky bucket, up to a maximum bucket size of 5 units; otherwise, one unit is subtracted. When the “water” in the leaky bucket reaches 5 units, the work space is declared to be occupied. Conversely, when the “water” in the bucket reaches 0, the work space is declared to be unoccupied. When SPOT thinks that the work space is occupied, it evaluates the current  $ppv(x)$  value and compares it to the desired comfort set point  $dc$ . At the beginning of each minute, if  $ppv(x) < dc$ , the heater is turned on until the PPV value reaches  $dc$ ; otherwise the heater is turned off. The detailed heater control logic is demonstrated in Figure 5.

## 5. EVALUATION

We discuss the evaluation of our system in this section. Since the evaluation period was in winter, we validated our design by implementing a heating control system. Our ideas, however, are applicable to a cooling system, where we would replace the heater with a small personal fan.

To evaluate the effectiveness of our system, we deployed the prototype to a  $11.9m^2$  office room at the University of Waterloo. The office room is owned by a professor and it is usually occupied from 8:30 AM to 5:30 PM on weekdays. Note that the room is in a building that also has its own HVAC control system whose design goal is to maintain a constant temperature of  $23^\circ C$  throughout the day. Therefore, the PPV setpoint is chosen to correspond to a comfort level that is somewhat warmer than usual (corresponding to a worker who prefers warm working conditions), as a positive offset to this nominal base value.

### 5.1 Accuracy of Clothing Level Estimation

We first discuss the effectiveness of the clothing level estimation. Since the system is designed for indoor thermal control, we assume that the clothing level is between 0.7 (a shirt) and 1.3 (shirt, sweater and jacket), which are common in our office environment. In the training phase, 23 data points were collected as training data, with the clothing level ranging from 0.7 to 1.25. This allowed us to




Figure 6: Clothing level estimated by human v.s. estimated by our algorithm. The RMSE of the estimation is 0.8466 and the Pearson correlation coefficient was 0.9201 indicating good linear correlation.

compute a linear regression to estimate the clothing level from the infrared sensor reading.

Subsequently, about 20 volunteers were selected to participate in a test of accuracy. For volunteers wearing a jacket, we first tested the clothing level estimation algorithm when they were wearing their jackets. We then tested the clothing level after they took off their jackets. Therefore, we collected 35 testing data points in total.

To test our algorithm, the clothing level of each volunteer was first estimated by one of the authors using Table 12. It was then evaluated using the estimation algorithm. The results are shown as a scatter plot in Figure 6. The root mean square error (RMSE) of the prediction was 0.8466 and the Pearson correlation coefficient was 0.9201 indicating good linear correlation.

Note that the infrared sensor we are using has a 5 degree detection angle. We found that when a subject was more than 2 meters away from the sensor, the clothing estimation result was inaccurate because of noise from background infrared radiation. Therefore, in a real deployment, we need to install the IR sensor no more than 2 meters from the worker. If this is an issue, for example in a large office, we advocate using sensors with a smaller detection angle.

### 5.2 Accuracy of PPV Estimation

This subsection discusses the accuracy of comfort level estimation using the PPV model. The evaluation was done in an actual office at the University of Waterloo for several days.

On the first day, the office owner gave votes to the system on the thermal environment to train the PPV model. Over the training period, 12 votes were collected as training data. We then tested the PPV model by comparing the predicted votes with 8 actual votes on the following days. The results are plotted in Figure 7. We found that the RMSE of PPV estimation was 0.5377 and the Pearson correlation coefficient was 0.8182 indicating good linear correlation.

### 5.3 Responsiveness of the Work Space to Thermal Control

This section discusses the responsiveness of the experimental workspace to thermal control using the radiant heater. To test re-






Figure 7: Actual personal vote v.s. predicted personal vote. The RMSE of the estimation is 0.5377 and the Pearson correlation coefficient was 0.8182 indicating good linear correlation.




Figure 8: Room temperature and the heater state. The room temperature increased 1 degree Celsius after turning on the heater for 15 minutes.

sponsiveness, we turned on the radiant heater at 4:20 PM and measured the PPV every minute. The result is plotted in Figure 8.

Before the heater was turned on at 4:20 PM, the room temperature was maintained by the central HVAC system and the PPV was between -0.5 to -0.4. When the heater was turned on at 4:20 PM, the room temperature as well as the PPV started to increase immediately, reaching the target PPV of 0 in 15 minutes at 4:35 PM.

We conclude that it is feasible to use reactive control for personal thermal comfort without significantly reducing human comfort.

## 5.4 Thermal Comfort Over a Day

We now discuss the performance of SPOT over the course of a typical day. Our experimental setup was the same as for the other experiments. However, we required SPOT to maintain a PPV of 0. Figure 9 shows the results of this experiment, depicting room occupancy, PPV, and room temperature over the day.

Note that both the room temperature and PPV are relatively low before the worker entered the office at 10 AM in the morning, with the PPV of -1.5 being the comfort level corresponding to the temperature setpoint chosen by the central HVAC system. SPOT turned on the heater within five minutes of the worker's arrival and both




Figure 9: Room occupancy and PPV over a day. Each tick on the X axis is 10 minutes. For most of the time when the room is occupied, the PPV is maintained around 0, even through there are external disturbances from the central HVAC system.

the temperature and the PPV increased steadily to 0 over the next 45 minutes.

The PPV of the office was always maintained around 0 when the worker was in the office, with small excursions above zero when the HVAC heating system turned on from time to time. The brief change in occupancy just before 11 AM, of about 10 minutes, was too short to cause any appreciable change in PPV.

The worker left the office at 4 PM for an hour. During that time, SPOT turned off the heater to save energy. This reduced the PPV to -0.5, but the PPV returned to 0 soon after the worker returned at 5 PM. When the worker finally left at 5:30 PM, the PPV declines, eventually reaching -1.5.

This demonstrates that SPOT can maintain the PPV at a chosen comfort value over the course of a day, despite the periodic activation of the central HVAC heating system and changes in office occupancy.

Note that, in this instance, PPV tracks room temperature quite closely. This is because there was little change in other environmental and personal factors, such as humidity and clothing level. In other circumstances, such as when the worker may put on or take off a jacket, SPOT would be able to maintain the comfort level by appropriately reducing the room temperature.

## 5.5 Trade-off between PPV and Energy Consumption

SPOT allows a building's energy consumption to be decreased in three ways.

- It allows the common areas of the building to be heated or cooled to a lesser degree than the ASHRAE standard of 23°C.
- It only heats or cools a work space when the worker is actually present.
- It allows the worker to choose a comfort level that is lower than 0, thus saving energy.

Here, we focus on the third element above.

To evaluate the possible amount of energy saving by lowering the PPV value, we measured the relationship between PPV and heater energy consumption<sup>5</sup>. We did this by setting the heater power to different values and recording the PPVs when the room temperature

<sup>5</sup>This relationship is necessarily noisy because temperature is not the only determinant of PPV. Nevertheless, the trend is distinct.




Figure 10: Relationship between daily energy consumption of the heater and the PPV. Maintaining a PPV of 0 consumes about 6 kWh electricity daily. By setting the target PPV to -0.25, we can save about 1.5 kWh electricity per day.

had converged. When the heater was turned off, the room temperature was maintained by the centralized HVAC system at around 23 degrees, corresponding to PPV values between -0.5 and -1.24 depending on the central HVAC system’s phase in its heating cycle. In contrast, when the radiant heater was set to its maximum power, the PPV was about 0.75 with the estimated power consumption per day was about 10.5 kWh.

Figure 10 shows this trade-off between PPV and the heater energy consumption in a day. We see that a reduction in PPV of 0.1, which is hardly noticeable by a human, results in the reduction in usage of 0.6 kWh of electricity in a day. This allows us to quantitatively select the trade-off between personal thermal comfort and heating energy consumption. For instance, an energy-aware office worker can set the target PPV value  $dc$  (as mentioned in §4.4) to -0.5 in order to save energy.

## 6. RELATED WORK

### 6.1 PMV Model

The PMV model is widely used for evaluating the performance of building temperature control systems. Yang et. al. [28] used PMV to reduce the energy consumption of a building’s HVAC system. Since radiant temperature is a significant factor in PMV in hot and humid areas, air temperature and humidity control is not enough for cooling in summer. In their system, they control the air velocity as well in order to maintain PMV at the comfortable range. Aswani [7] et. al. have also used the PMV model in their building temperature control system. They use Learning-Based Model Predictive Control (LBMPC) to control the building HVAC system such that different zones of the building maintain PMVs close to 0. Our approach, instead, uses the personalized PPV model to achieve personalized thermal control.

The Thermovote [12] system allows workers in a building to vote on the current temperature. Instead of using PMV to predict users’ feeling, they use the actual vote of the workers to adjust their comfort. However, the system requires the users to vote frequently, which is onerous. We avoid this problem by building personalized models for each individual. SPOT only requires votes during the training phase to calibrate the PPV model. Subsequently, the thermal preference of the user is used to control the HVAC system and no more voting is required.

### 6.2 Occupancy-based HVAC Control

There has been a considerable amount of work on improving the energy-efficiency of HVAC systems. For example, Aswani et. al. [7] use Learning-Based Model Predictive Control (LBMPC) to model and control HVAC systems in a large university building. They were able to save an average of 1.5MWh of electricity per day in their testbed. Fong et. al [16] used evolutionary programming (EP) to find the optimal HVAC setting, and apply this setting instead of the default one. They found that about 7% of energy could be saved by replacing the default HVAC settings by their optimized ones. Although sophisticated, these approaches control only the room temperature, rather than attempt to achieve a certain level of user comfort, as we do.

Turning off an HVAC system when no humans are present is an obvious technique to reduce energy use. It is also important to turn on heating in advance of human occupancy, because it can take tens of minutes to heat a cold building to tolerable levels. Most occupancy prediction methods use previously collected occupancy data. Lu et. al. [21] showed that by learning occupancy and sleep patterns, it is possible to save about 28% of energy use in a home environment. The PreHeat system [27] uses occupancy sensors to predict home occupancy patterns and automatically adjusts the HVAC temperature setpoint to save energy. If we consider the previous history of work space occupancy as a time series, and current time as a function of previous entries in the time series, we can learn this function as a Gaussian Process. This approach has been used in Erickson’s and Rogers’ papers [14, 26]. In another learning-based approach, home occupancy is modeled as a Markov chain and room occupancy is encoded as a state in the Markov model [13]. Mozer et. al. use a neural network and a lookup table to predict home occupancy in their Neuralthermostat project [23]. To build human interpretable model, Leephakpreeda [19] applied a grey model for occupancy prediction. Ardakanian et. al [8] uses sound and light level of a room to infer occupancy and applies POMDP for optimal HVAC control.

Most learning based models require relatively large amount of data in order to produce accurate predictions. Hence, to predict occupancy with limited historical occupancy data, there exists other approaches to employ some side channels to assist prediction. For example, Gupta et. al. [17] uses GPS sensors on mobile phones to estimate the arrival time of home owners and heat the house before they arrive.

HVAC control based on occupancy can be used in conjunction with our techniques to allow the HVAC controller to pre-heat or pre-cool a work space to achieve a target comfort level rather than a target temperature. We discuss this further in §7.3

## 7. DISCUSSION AND FUTURE WORK

### 7.1 Extreme Sensing

The SPOT system, with its plethora of sensors, can be viewed as a somewhat extremal point in the space of HVAC control systems. We are keenly aware that our approach is hardware and compute intensive, and has a price point that may put it out of reach of most offices. Nevertheless, we believe that our approach is interesting for at least two reasons.

First, with the proliferation of sensing and compute systems, even high-end sensors such as the Microsoft Kinect will be much cheaper in the near future. Second, even if maintaining per-worker comfort is too expensive in terms of sensing, per-worker temperature control, a far more achievable goal, is cheap, effective, and well within reach in existing offices. We believe, therefore, that SPOT establishes an interesting data point in the thermal control design space.

## 7.2 Predictive Control

Our system uses reactive control to adjust the temperature of the room. An alternative approach is to use Model Predictive Control (MPC) [9] to adjust the setpoint of the HVAC system. MPC is fundamentally different from reactive control because it builds a thermal model for the system. Since we can estimate the internal mechanics of the system, we can predict the future control outputs given the current control inputs and the states of the system. Model Predictive Control is a white box approach, hence it is usually easier to tune the model parameters to meet the optimal control objective. For example, by using MPC, we can build a thermal model of the room using the heater power as the input. With this model, we can easily find the optimal control strategy that minimizes the energy consumption (i.e., the integral of power). We decided not to use this more complex approach because simple reactive control appears to be adequate for a small work space with negligible thermal mass.

## 7.3 Incorporating Occupancy Prediction

We have already discussed many well-known approaches for occupancy prediction. Our system can take advantage of existing occupancy prediction algorithms to make control decisions *before* room occupancy changes. For example, in Figure 9, in the morning, it took about 45 minutes to get the PPV to 0. Instead, with occupancy prediction, as in Preheat [27], we could start to heat the room 1 hour before the estimated arrival time of the worker and maintain a PPV of -0.5. We could then heat the room to the target PPV of 0 only when the worker actually arrived.

Similarly, the thermal mass of a work space causes it to cool down over many tens of minutes. Therefore, we can stop heating the work space in advance of the worker leaving it by using an occupancy prediction algorithm [11]. We intend to explore this direction in future work.

Note that prediction accuracy critically affects the performance control system. False negative prediction will reduce worker comfort and false positive prediction will lead to a waste of energy. Therefore, when using occupancy prediction, we must evaluate prediction accuracy and the cost of false predictions.

## 7.4 Optimal Control

With accurate occupancy prediction, we can use an optimal control framework to further reduce energy consumption. For example, in Figure 9, when the heater was turned off at 5:30 PM, the room temperature was at 26° C. It took 2.5 hours for the office to cool down to 24° C. If the office had very good insulation, the cool down process would have been even slower. Therefore, we can save energy by turning off the heater earlier if we know that the office will be unoccupied in the near future. Conversely, if we can predict the arrival time of the worker, we can heat the room up to a comfortable temperature before the worker arrives. However, to take advantage of these approaches, we need to decide how much time in advance should we turn on or turn off the heater. An optimal control framework would allow us to decide the best timing for any control action over a planning horizon.

## 7.5 Human Factors in Automation

Our discussion so far has assumed that the worker has little role to play in thermal control. In fact, workers themselves can be active participants in a thermal control system if they receive and act on energy-saving tips. For example, SPOT could, instead of turning on a heater, suggest to workers that they put on a jacket. This integration of humans into the control loop can be viewed as being unnecessarily intrusive. Nevertheless, we believe that, if properly

presented to humans, such control actions can be both energy saving and marginally intrusive. We intend to explore this in future work.

## 7.6 Limitations

Our work has several inherent limitations that we discuss next.

**Thermal isolation** We assume that each work space is relatively thermally isolated from other work spaces. This does not hold true, for example, in open-plan offices.

**Personalized work spaces** We assume that each worker has their own personal space, and that they do not move from space to space over time. This may not be a valid assumption for all workplaces.

**Cost** Each SPOT system costs about \$1,000. This may be too high a cost to pay for modest increases in worker comfort. We expect this cost to rapidly decline over the next few years.

**Calibration** SPOT requires worker participation to calibrate personal comfort levels. This can take a day or so, and can be viewed as onerous by some workers.

**Validation** Because we do not have control over our building's temperature setpoint, we are unable to validate our research hypothesis that personalized thermal control can save energy overall. This is certainly quite plausible, in that measurements have shown that a two degree increase in the temperature setpoint in summer can reduce home energy use by 37% [1], but we have no way to validate this conclusion.

**Environment** SPOT is blind to windows that are open versus closed, to HVAC state, and user mobility. In our experiment, the office is controlled by a centralized heating system and the window is always closed. These factors may affect the effectiveness of SPOT in other environment.

## 8. CONCLUSION

We have presented the design and implementation of SPOT, a smart personal thermal comfort system for office work spaces. SPOT builds on three underlying ideas. First, we extend the PMV model to create the PPV model to quantitatively estimate personal comfort. Second, we use a set of sensors, including a Microsoft Kinect sensor, to measure PPV model parameters, so that we can estimate a worker's comfort level at any point in time. Third, we use occupancy-aware simple reactive control to maintain the PPV despite changes in the environment. We deployed SPOT in a real office environment and validated that it can maintain comfort over the course of a typical work day. Moreover, we have shown how SPOT allows a worker to trade off a reduction in comfort for saving energy. We believe that our work demonstrates an interesting case study of how to maintain human comfort using extreme sensing. Finally, a limited version of our system, that only maintains personalized temperature offsets from a building-wide base setpoint, is not only easy to deploy, but is also likely to reduce overall building energy use.

## 9. REFERENCES

- [1] 2009 peaksaver Residential Air Conditioner Measurement and Verification Study. *Ontario Power Authority, May 2010.*
- [2] ANSI/ASHRAE Standard 55-1992, Thermal Environmental Conditions for Human Occupancy. *Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers.*



- [3] Energy Efficiency Fact Sheet: Heating, ventilation & Air Conditioning. [http://www.originenergy.com.au/files/SMEfs\\_HeatingAirCon.pdf](http://www.originenergy.com.au/files/SMEfs_HeatingAirCon.pdf).
- [4] Ergonomics of the Thermal Environment - Analytical Determination and Interpretation of Thermal Comfort using Calculation of the PMV and PPD Indices and Local Thermal Comfort Criteria. *ISO 7730:2005*.
- [5] HVAC & Energy Systems. <http://canmetenergy.nrcan.gc.ca/buildings-communities/hvac-energy/908>.
- [6] Nest, the learning thermostat. <http://www.nest.com/>.
- [7] Anil Aswani and Neal Master and Jay Taneja and Andrew Krioukov and David E. Culler and Claire Tomlin. Energy-Efficient Building HVAC Control Using Hybrid System LBMPC. *CoRR*, abs/1204.4717, 2012.
- [8] O. Ardakanian and S. Keshav. Using Decision Making to Improve Energy Efficiency of Buildings. *ICAPS-10 POMDP Practitioners Workshop, May 2010*.
- [9] E. Camacho, C. Bordons, E. Camacho, and C. Bordons. *Model Predictive Control*, volume 303. Springer Berlin, 1999.
- [10] R. de Dear and G. S. Brager. Developing an Adaptive Model of Thermal Comfort and Preference. *UC Berkeley: Center for the Built Environment*, 1998.
- [11] C. Ellis, J. Scott, M. Hazas, and J. Krumm. EarlyOff: Using House Cooling Rates to Save Energy. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*.
- [12] V. Erickson and A. Cerpa. Thermovote: Participatory Sensing for Efficient Building HVAC Conditioning. 2012.
- [13] V. L. Erickson and A. E. Cerpa. Occupancy based Demand Response HVAC Control Strategy. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, BuildSys '10, pages 7–12, New York, NY, USA, 2010. ACM.
- [14] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan. Energy efficient building environment control strategies using real-time occupancy measurements. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '09, pages 19–24, New York, NY, USA, 2009. ACM.
- [15] P. O. Fanger. Thermal comfort. *Analysis and Applications in Environmental Engineering*, Danish Technical Press, Copenhagen, Denmark, 1970.
- [16] K. Fong, V. Hanby, and T. Chow. HVAC System Optimization for Energy Management by Evolutionary Programming. *Energy and Buildings*, 38(3):220–231, 2006.
- [17] M. Gupta, S. S. Intille, and K. Larson. Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, pages 95–114, Berlin, Heidelberg, 2009. Springer-Verlag.
- [18] B. W. Jones. Capabilities and Limitations of Thermal Models for Use in Thermal Comfort Standards. *Energy and Buildings*, 34(6):653–659, 2002. Special Issue on Thermal Comfort Standards.
- [19] T. Leephakpreeda. Adaptive Occupancy-based Lighting Control via Grey Prediction. *Building and Environment*, 40(7):881–886, 2005.
- [20] W. Liping and W. N. Hien. The Impacts of Ventilation Strategies and Facade on Indoor Thermal Environment for Naturally Ventilated Residential Buildings in Singapore. *Building and Environment*, 42(12):4006–4015, 2007.
- [21] 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*, SenSys '10, pages 211–224, New York, NY, USA, 2010. ACM.
- [22] G. K. C. Mehdi Shahrestani, Runming Yao. Performance Characterisation of HVAC&R Systems for Building Energy Benchmark. <http://www.cibse.org/content/cibsesymposium2012/Poster026.pdf>.
- [23] M. Mozer, L. Vidmar, R. Dodier, et al. The Neurothermostat: Predictive Optimal Control of Residential Heating Systems. *Advances in Neural Information Processing Systems*, pages 953–959, 1997.
- [24] J. Nicol and M. Humphreys. Adaptive Thermal Comfort and Sustainable Thermal Standards for Buildings. *Energy and Buildings*, 34(6):563–572, 2002. Special Issue on Thermal Comfort Standards.
- [25] L. Prez-Lombard, J. Ortiz, and C. Pout. A Review on Buildings Energy Consumption Information. *Energy and Buildings*, 40(3):394–398, 2008.
- [26] A. Rogers, S. Maleki, S. Ghosh, and J. Nicholas R. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *Second International Workshop on Agent Technology for Energy Systems (ATES 2011)*, pages 71–78, May 2011. Event Dates: May 2011.
- [27] J. Scott, A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. PreHeat: Controlling Home Heating Using Occupancy Prediction. In *Proceedings of the 13th international conference on Ubiquitous computing*, UbiComp '11, pages 281–290, New York, NY, USA, 2011. ACM.
- [28] K. Yang and C. Su. An Approach to Building Energy Savings Using the PMV Index. *Building and Environment*, 32(1):25–30, 1997.

## 10. APPENDIX: DETAILS OF THE PMV MODEL

The PMV [15] is computed as.

$$pmv(\mathbf{x}) = (0.303 \cdot \exp(-0.036 \cdot M) + 0.028) \cdot \left\{ \begin{array}{l} (M - W) - 3.05 \cdot 10^{-3} \cdot (5733 - 6.99 \cdot (M - W) - p_a) \\ -0.42 \cdot ((M - W) - 58.15) - 1.7 \cdot 10^{-5} \cdot M \cdot (5867 - p_a) \\ -0.0014 \cdot M \cdot (34 - t_a) - 3.96 \cdot 10^{-8} \cdot f_{cl} \cdot ((t_{cl} + 273)^4 \\ - (\bar{t}_r + 273)^4) - f_{cl} \cdot h_c \cdot (t_{cl} - t_a) \end{array} \right\} \quad (7)$$

where  $t_{cl}$  is the clothing surface temperature, and  $W$  is the effective mechanical power which is 0 for most indoor activities.

Variable  $t_{cl}$  can be evaluated by:

$$t_{cl} = 35.7 - 0.028 \cdot (M - W) - I_{cl} \cdot (3.96 \cdot 10^{-8} \cdot f_{cl} \cdot ((t_{cl} + 273)^4 - (\bar{t}_r + 273)^4) + f_{cl} \cdot h_c \cdot (t_{cl} - t_a)) \quad (8)$$

Variable  $h_c$  is the convective heat transfer coefficient, which is derived as

$$h_c = \begin{cases} 2.38 \cdot |t_{cl} - t_a|^{0.25} & \text{if } 2.38 \cdot |t_{cl} - t_a|^{0.25} > 12.1 \cdot \sqrt{v_{ar}} \\ 12.1 \cdot \sqrt{v_{ar}} & \text{if } 2.38 \cdot |t_{cl} - t_a|^{0.25} < 12.1 \cdot \sqrt{v_{ar}} \end{cases} \quad (9)$$

Variable  $f_{cl}$  is the clothing surface area factor, which is derived as:

$$f_{cl} = \begin{cases} 1.00 + 1.290I_{cl} & \text{if } I_{cl} \leq 0.078m^2 \cdot K/W \\ 1.05 + 0.645I_{cl} & \text{if } I_{cl} > 0.078m^2 \cdot K/W \end{cases} \quad (10)$$

In practice, the metabolic rate and the clothing insulation are first estimated by Table 11 and Table 12. Given the clothing insulation  $I_{cl}$ , we calculate the clothing surface temperature  $t_{cl}$  and the convective heat transfer coefficient  $h_c$  by iteratively applying Equation 8 and 9. Finally, by using Equation 7 and 10, we can estimated the Predicted Mean Vote.

| Activity                  | Metabolic Rate |       |
|---------------------------|----------------|-------|
|                           | $W/m^2$        | $met$ |
| Reclining                 | 46             | 0.8   |
| Seated, relaxed           | 58             | 1.0   |
| Sedentary activity        | 70             | 1.2   |
| Standing, medium activity | 93             | 1.6   |

Table 11: Metabolic Rates

| Daily Wear Clothing   | Clothing Insulation ( $I_{cl}$ ) |                 |
|---|----------------------------------|-----------------|
|   | $clo$                            | $m^2 \cdot K/W$ |
| Panties, T-shirt, shorts, light socks, sandals  | 0.30                             | 0.050           |
| Underpants, shirt with short sleeves, light trousers, light socks, shoes                    | 0.50                             | 0.080           |
| Panties, petticoat, stockings, dress, shoes   | 0.70                             | 0.105           |
| Underwear, shirt, trousers, socks, shoes  | 0.70                             | 0.110           |
| Panties, shirt, trousers, jacket, socks, shoes  | 1.00                             | 0.155           |
| Panties, stockings, blouse, long skirt, jacket, shoes                                       | 1.10                             | 0.170           |
| Underwear with long sleeves and legs, shirt, trousers, V-neck sweater, jacket, socks, shoes | 1.30                             | 0.200           |
| Underwear with short sleeves and legs, shirt, trousers, vest, jacket, coat, socks, shoes    | 1.50                             | 0.230           |

Table 12: Thermal Insulation for different clothing level