Optimal Personal Comfort Management Using SPOT+

Peter Xiang Gao
David R. Cheriton School of Computer Science
University of Waterloo
peterxianggao@gmail.com

S. Keshav
David R. Cheriton School of Computer Science
University of Waterloo
keshav@uwaterloo.ca

ABSTRACT

We present SPOT+, a system that allows office workers to optimally balance between heating energy consumption and personal thermal comfort. In prior work, we described SPOT: a smart personal thermal control system based on *reactive* control [8]. In contrast, the SPOT+ system performs *predictive* control. Specifically, SPOT+ uses the k-nearest-neighbour algorithm to predict room occupancy and learning-based model predictive control (LBMPC) to predict future room temperature and to compute the optimal sequence of control inputs. This allows the system to schedule future temperature setpoints to optimize an objective function expressed as a linear combination of thermal comfort and energy consumption. We have deployed SPOT+ as well as four other alternative control schemes in an office workspace. We find that SPOT+ reduces energy usage by 60% compared to a fixed-temperature setpoint and reduces personal thermal discomfort from 0.36 to 0.02 (in the ASHRAE comfort scale) compared to SPOT.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Design, Performance

Keywords

Energy management

1. INTRODUCTION

Buildings account for almost a third of global energy consumption [11]. Heating, Ventilation and Air-Conditioning (HVAC) systems are the dominant energy consumers in both residential and commercial buildings in most developed countries, accounting for 50%-70% of their energy use [1,2,13,15]. Reducing the energy consumption of HVAC systems can, therefore, significantly reduce energy consumption in these countries [17].

We focus on heating individual workspaces[1] during the winter season. The energy consumption of the workspace heating system obviously depends on the chosen temperature setpoint: the higher this setpoint, the greater the energy use. Thus, it is always possible to save energy either by lowering the setpoint or by turning the system off when a workspace is presumed vacant. However, both actions reduce worker thermal comfort, so overly aggressive temperature setback is counter-productive. Thus, our overall goal is to allow individual workers to balance between energy use and their own comfort.

In prior work [8], we described SPOT, a personal thermal control system for workspaces that models worker comfort using the Predicted Personal Vote (PPV) model (described below) and automatically adjusts room heating to maintain a desired comfort level. A building using SPOT could lower the overall building temperature setpoint, with a SPOT controller in each work space providing 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 heats the work space to a personalized higher level. Importantly, SPOT is *reactive* in that it only takes control actions when the worker is actually present. Therefore, it cannot pre-heat the workspace before the arrival of the worker, or turn off heating in anticipation of the worker's departure. In this paper, we present SPOT+ that *learns* worker occupancy patterns and room thermal characteristics to compute the optimal sequence of control inputs.

Specifically, SPOT+ uses the k-nearest-neighbour algorithm to predict future workspace occupancy from past observations [16]. It also builds a thermal model of the room using learning-based model predictive control (LBMPC) so that it can predict the future workspace temperature given the energy input of the HVAC system. It then uses an optimal control framework to find the best temperature setpoint schedule for the rest of the day, over a set of half-hour timeslots. Such an optimal setpoint schedule preheats the room before the estimated worker arrival time and stops heating before the estimated worker departure time. It also finds the optimal setpoint that minimizes energy consumption without affecting user comfort when the room is occupied.

[1]For simplicity, we assume that workspaces are thermally isolated. Thus, our work does not apply to open-plan offices.

BuildSys '13 November 14-15, Roma, Italy. Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-2431-1/13/11.$15.00. http://dx.doi.org/10.1145/2528282.2528297

Our work makes the following contributions:

- We have designed SPOT+, an optimal predictive control scheme for personalized workspace thermal control that balances energy consumption and worker comfort
- We have implemented SPOT+ as well as four other temperature control schemes and compared their relative performance in a real testbed
- We find that SPOT+ can save about 60% of energy comparing with a fixed temperature setpoint, and can reduce thermal discomfort from 0.36 to 0.02 (on the ASHRAE comfort scale) compared to the SPOT reactive temperature control scheme

The rest of this paper is laid out as follows. Section II presents a background on quantitative comfort modelling and an overview of SPOT. We discuss optimal control using LBMPC as well as occupancy prediction in Section III, followed by an evaluation of SPOT+ and competing schemes in Section IV. Section V presents related work and we conclude in Section VI.

2. BACKGROUND

We first describe the Predicted Personal Vote (PPV) comfort metric to evaluate personal thermal comfort in indoor environments that we used in the SPOT control system [8]. We then describe SPOT and it's shortcomings, which motivate SPOT+.

2.1 Predicted Personal Vote (PPV) Model

The PPV model is a generalization of the well-known Predicted Mean Vote (PMV) model [3, 7]. The PMV model estimates an average worker's comfort level on the 7-point ASHRAE scale[2] using a function $f_{pmv}(\cdot)$:

$$pmv = f_{pmv}(\mathbf{x}) = f_{pmv}(t_a, \bar{t_r}, v_{ar}, p_a, M, I_{cl}) \tag{1}$$

where pmv is the predicted mean vote and $\mathbf{x}$ denotes the following environmental and personal variables:

- 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 worker
- I_{cl} is the worker's clothing insulation factor

Then,

$$pmv(\mathbf{x}) = (0.303 \cdot exp(-0.036 \cdot M) + 0.028) \cdot \left\{ \begin{array}{c} (M-W) - 3.05 \cdot 10^{-3} \cdot (5733 - 6.99 \cdot (M-W) - p_a) \\ -0.42 \cdot ((M-W) - 58.15) - 1.7 \cdot 10^{-5} \cdot M \cdot (5867 - p_a) \\ -0.0014 \cdot M \cdot (34 - t_a) - 3.96 \cdot 10^{-8} \cdot f_{cl} \cdot ((t_{cl} + 273)^4 \\ -(\bar{t_r} + 273)^4) - f_{cl} \cdot h_c \cdot (t_{cl} - t_a) \end{array} \right\} \tag{2}$$

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

[2]Cold (-3), Cool (-2), Slightly Cool (-1), Neutral (0), Slightly Warm (+1), Warm (+2), and Hot (+3).

Variable t_{cl} can be evaluated by:

$$\begin{aligned} t_{cl} =& 35.7 - 0.028 \cdot (M-W) - I_{cl} \cdot (3.96 \cdot 10^{-8} \cdot f_{cl} \cdot \\ & ((t_{cl} + 273)^4 - (\bar{t_r} + 273)^4) + f_{cl} \cdot h_c \cdot (t_{cl} - t_a)) \end{aligned} \tag{3}$$

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

$$h_c = \begin{cases} 2.38 \cdot |t_{cl} - t_a|^{0.25} & \text{if } 2.38 \cdot |t_{cl} - t_a|^{0.25} > 12.1 \cdot \sqrt{v_{ar}} \\ 12.1 \cdot \sqrt{v_{ar}} & \text{if } 2.38 \cdot |t_{cl} - t_a|^{0.25} < 12.1 \cdot \sqrt{v_{ar}} \end{cases} \tag{4}$$

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

$$f_{cl} = \begin{cases} 1.00 + 1.290 I_{cl} & \text{if } I_{cl} \le 0.078 m^2 \cdot K/W \\ 1.05 + 0.645 I_{cl} & \text{if } I_{cl} > 0.078 m^2 \cdot K/W \end{cases} \tag{5}$$

When the PPV of an office space is estimated by a human expert, the office worker's metabolic rate and level of clothing insulation are first estimated using Table 2 and Table 3 in the Appendix. Given the clothing insulation I_{cl}, it is possible to calculate the clothing surface temperature t_{cl} and the convective heat transfer coefficient h_c by iteratively applying Equation 3 and 4. Finally, by using Equation 2 and 5, the Predicted Mean Vote can be estimated.

Although this seminal model was developed in 1970, recent work has validated its accuracy for climate-controlled environments based on field studies in 160 buildings located in varied climatic zones [6].

To evaluate comfort in workspaces occupied by a *single* office owner, SPOT computes a Predicted Personal Vote (PPV) as an affine transform of pmv:

$$ppv = f_{ppv}(pmv) \tag{6}$$

This function is learnt using least squares linear regression during a training phase, with the worker providing ground truth on comfort level.

2.2 SPOT

The key idea behind SPOT is to control comfort in a *single* workspace by sensing the six variables underlying the PPV model in Eq. 1. A standard environment sensor is used to measure air temperature, radiant temperature, air velocity and humidity. Given that this is an office environment, the metabolic rate can be assumed to be constant and low. Therefore, the most difficult variable to sense is the clothing level.

Figure 1 shows the SPOT clothing sensor built using a Microsoft Kinect and a servo-controlled infrared sensor. The Kinect tracks worker location; this information is fed to the servos that adjust their rotation angle to point an infrared sensor to the worker's chest. Knowing that the infrared radiation emitted by human body is negatively correlated to the clothing level, SPOT solves an inverse problem using linear regression to estimate the clothing level as described in [8]. This work also shows that SPOT accurately estimates the PMV (and PPV) despite changes in clothing levels and uncontrolled changes to room temperature by the building's heating system.

2.3 SPOT shortcomings

The Kinect sensor allows SPOT to be reasonably sure about the true occupancy status of the workspace, but it

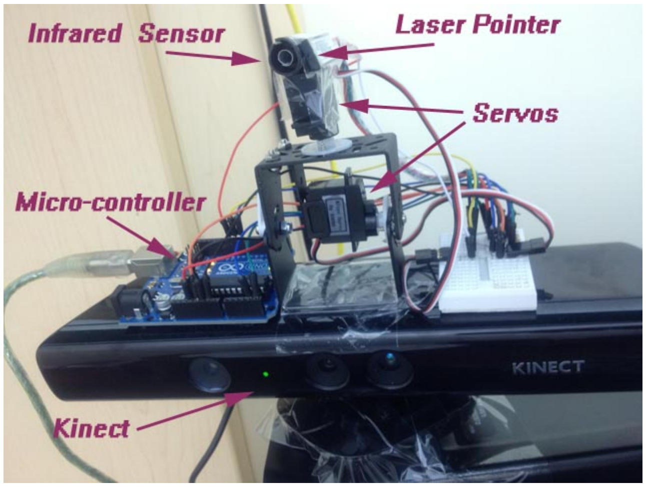

Figure 1: SPOT is built using a Microsoft Kinect sensor, an Arduino Microcontroller, and an infrared sensor. The Kinect tracks the location of the user. The servos control the rotation angle of the infrared sensor to make sure that it always points to the torso of the worker. The infrared sensor detects the infrared radiation emitted by human body. By measuring the infrared attenuation by clothing, SPOT can estimate the clothing level of the worker. The laser pointer is used to calibrate the servos.

only *reacts* to occupancy. That is, it does not heat the workspace until the worker is actually present. However, it takes some time for the workspace to warm up, so the worker may feel uncomfortable when first entering it on a cold morning. Moreover, the workspace temperature may continue to be high even after it becomes vacant: it would be more efficient to let the temperature drop slightly just before the worker departs. This motivates us to exploit ground-truth knowledge of worker occupancy to predict future occupancy, thus allowing us to improve worker comfort and reduce energy cost.

3. CONTROLLER DESIGN

3.1 Learning-Based Model Predictive Control (LBMPC)

Traditional feedback control determines the control input based on the control error, i.e., the gap between the desired and the actual output. For example, if the room temperature is lower then the desired setpoint, the heater (control input) is turned on. LBMPC adds two refinements.

- First, machine learning is used to learn a *model* of the physical system.
- Second, the model is used to choose *optimal* control inputs such that a future control goal is met.

Note that LBMPC does not require an error term in order to take control actions, thus providing the benefits of open-loop control.

SPOT+ learns a *thermal model* of the workspace. An example of a fact that can be derived from such a model is "if the heater power is set to 800W, the workspace temperature increases from 22°C to 23°C in 20 minutes." This model is used to determine the appropriate heater power to increase the room temperature to a certain point.

We now describe the mathematical formulation of room thermal model. Given the outside temperature T_{out} and the indoor temperature T_{in}, by the Newton's Law of Cooling, the rate of thermal energy loss, P_{loss} is proportional to the temperature difference:

$$P_{loss} = k(T_{in} - T_{out}) \tag{7}$$

where k is the conduction factor of the workspace: a workspace with better insulation has a smaller conduction factor. Suppose that workspace heating is achieved by a heater with power P_{hvac} and efficiency e. The net heat input rate, P, is given by:

$$P = eP_{hvac} - P_{loss} = eP_{hvac} - k(T_{in} - T_{out}) \tag{8}$$

This net heat input rate is the derivative of thermal flux Q and is proportional to the temperature change:

$$P = \frac{dQ}{dt} = C\frac{dT_{in}}{dt} \tag{9}$$

where C is the heat capacity of the workspace. Combining Eq. 8 and Eq. 9, we have:

$$\frac{dT_{in}}{dt} = \frac{eP_{hvac} - k(T_{in} - T_{out})}{C} \tag{10}$$

To enable digital control, we convert Eq. 10 to its discrete version:

$$T_{in}(s+1) = T_{in}(s) + \frac{eP_{hvac}(s) - k(T_{in}(s) - T_{out}(s))}{C} \tag{11}$$

where $T_{in}(s)$ is the temperature at the s-th timestep.

The model contains three parameters: the efficiency of the HVAC system e, the conduction factor k, and the house heat capacity C. Given tuples of $\{T_{in}(s), T_{out}(s), P_{hvac}(s), T_{in}(s+1)\}$ at different timesteps, these model parameters can be estimated by linear regression. Thus, in the training phase, SPOT+ conducts a few controlled experiments when the room is unoccupied, then uses least square regression to find the parameters of thermal model.

3.2 Occupancy Prediction

It takes time to heat a cold workspace. If workspace occupancy can be predicted, we can use LBMPC to pre-heat it and increase worker comfort. Symmetrically, heating can be avoided if the workspace will not be occupied shortly thereafter (such as during a lunch break or at the end of a day). Inspired by the Pre-Heat system [16], we predict future occupancy at any time using the most similar occupancy history for that hour of the day. The advantage of this approach over most prior approaches is that it automatically corrects for changes in occupancy patterns during holidays and vacations, but has the problem that it performs poorly for the first part of the day.

Specifically, we maintain a database S of worker occupancy. We divide each day into 48 timeslots, each 30 minutes long. Let t be the identifier of a timeslot and let $TOD(t)$ return the index of the timeslot in the day (an integer in the range [0, 47]). For example, $TOD(t) = 3$ refers to the timeslot that starts at 1 AM and ends at 1:30 AM. Let $m(t)$ be the observed or predicted occupancy at timeslot t that is set to 1 if the room is occupied and 0 otherwise. Let the current time be t_0. To make a prediction of future occupancy, we first find all timeslots $s \in S$ such that $TOD(s) = TOD(t_0)$ and then compare their *similarity*, where similarity is defined as :

$$similarity(s,t) = \sum_{b=0}^{i} \delta_{s-b,t-b} \tag{12}$$

where $\delta_{s-b,t-b} = 1$ if $m(s-b) = m(t-b)$ and 0 otherwise. We select the top K timeslots with highest similarity to t_0 in S and we denote them as $s_k \in S_K$. To do j-th step ahead occupancy prediction, we calculate the occupancy probability using the following formula:

$$p(t+j) = \frac{1}{K}\sum_{k=1}^{K} m(s_k + j) \tag{13}$$

If the occupancy probability $p(t+j)$ is larger than a threshold, currently set to 0.5, we predict that timeslot $t+j$ will be occupied, i.e., $m(t+j) = 1$.

For our evaluation, we chose K to be 5 and we collected occupancy data for more than three months.

3.3 Optimal Control Strategy

At a high level, the goal of the optimal control algorithm is to decide the best time to turn on the heater, potentially before the predicted arrival time of the worker, and the best time to turn off the heater, potentially before the worker is predicted to leave. The control algorithm also provides a trade-off between worker comfort and energy saving.

More specifically, the optimal control strategy determines a heater operation sequence over an *optimization horizon* of H timesteps that minimizes a linear combination of the total energy use and a penalty term if the PPV does not lie in the acceptable range of of [-ϵ, ϵ] when the workspace is occupied[3]. Similar to Eq. 1, we let $\mathbf{x}(s)$ be the environmental and personal variables at time step s[4]:

$$\mathbf{x}(s) = \{t_a(s), \bar{t}_r(s), v_{ar}(s), p_a(s), M(s), I_{cl}(s)\}$$

where $t_a(s)$ is the air temperature at time step s. Function

$$ppv(\mathbf{x}(s))$$

evaluate the predicted personal vote (PPV) at time step s. Let $\beta_c(s)$ and $\beta_h(s)$ be the values of the cold and hot soft penalty terms respectively at time s, that is, the additional range of comfort that can be used, with a corresponding penalty, if this reduces energy use. Let λ be the relative weight given to thermal comfort. Then, the optimal control sequence is obtained by solving the following problem:

$$min \sum_{s=1}^{H} P_{hvac}(s) + \lambda \sum_{s=1}^{H} m(s)(\beta_c(s) + \beta_h(s)) \tag{14}$$

that minimize linear combination of energy use P_{hvac} and the discomfort penalty terms $\beta_c(s) + \beta_h(s)$. At every time step s, the problem is subject to the following constraints

$$\begin{aligned} ppv(\mathbf{x}(s)) &\geq -\epsilon - \beta_c(s) \\ ppv(\mathbf{x}(s)) &\leq \epsilon + \beta_h(s) \\ \beta_c(s) &\geq 0 \\ \beta_h(s) &\geq 0 \end{aligned}$$

The inequations give soft penalties to time steps that have absolute PPV values larger than ϵ.

Solving this optimization problem is difficult because the $ppv(.)$ function is non-linear and can be evaluated only using an iterative numerical method. Therefore, we convert the problem to a graphical problem where each graph node represents a potential system state over a ten-minute interval, each link represents a feasible state transition, and a link weight represents the energy and comfort cost of making a state transition. Then, finding the optimal control sequence corresponds simply to finding the shortest path in this graph.

Figure 2 shows an example of a graphical state model. Note that all the nodes at a particular time step form a *layer*. The r-th node at layer s, denoted $N_{s,r}$ represents a potential system state $\mathbf{x}_{s,r}$. For simplicity, we assume that the hot and cold penalty terms are identical and denoted $\beta(s)$ so that the penalty in the rth state at time step s, denoted $\beta(s,r)$ is given by:

$$\beta(s,r) = \begin{cases} 0, \; if \; |ppv(\mathbf{x}_{s,r})| < \epsilon \\ |ppv(\mathbf{x}_{s,r})| - \epsilon, \; otherwise \end{cases} \tag{15}$$

Edges connect all feasible state transitions between adjacent layers. The energy cost to transition from state $N_{s,r}$ to state $N_{s+1,r'}$ is denoted $P_{hvac}(N_{s,r}, N_{s+1,r'})$, which can be calculated from Eq. 11. The edge weight between nodes $N_{s,r}$ and $N_{s+1,r'}$ is the energy cost from state $N_{s,r}$ to state $N_{s+1,r'}$ plus the weighted comfort penalty:

$$d(N_{s,r}, N_{s+1,r'}) = P_{hvac}(N_{s,r}, N_{s+1,r'}) + \lambda m(s+1)\beta(s+1,r') \tag{16}$$

For example, if $\lambda = 10000$ and $m(2) = 1$, the edge weight between nodes $N_{1,1}$ and $N_{2,1}$ is $d(N_{1,1}, N_{2,1}) = 800 + 10000 * 1 * 0.5 = 5800$. The optimal control sequence is the shortest path from node $N_{1,1}$ to the dummy end node after the last step.

Note that the optimal control sequence updated at the end of each time step because of changes in the environment or if the predicted occupancy $m(s)$ is different from the actual observation.

4. EVALUATION

This section reports on a preliminary performance evaluation of five different temperature control schemes as well as the potential benefits of predictive over reactive control. The control target is a single office room in the University of Waterloo that was occupied by one of the authors[5]. The workspace is about $11.9m^2$ and its temperature is maintained at around 23°C by the centralized HVAC system. The worker usually comes to the office at around 8:30 AM and leaves at about 5:30 PM every weekday.

[3]In our evaluation, we choose H= 6, to obtain the optimal control sequence for the next hour.

[4]Note that $t_a(s)$ in Eq. 1 and $T_{in}(s)$ in Eq. 11 both denote the indoor air temperature at time step s.

[5]We realize that it would have been better for the evaluation to have been done with office worker unrelated to the authors, but, given the experimental nature of this work, we were reluctant to solicit volunteers who might freeze or sweat due to bugs in our software!

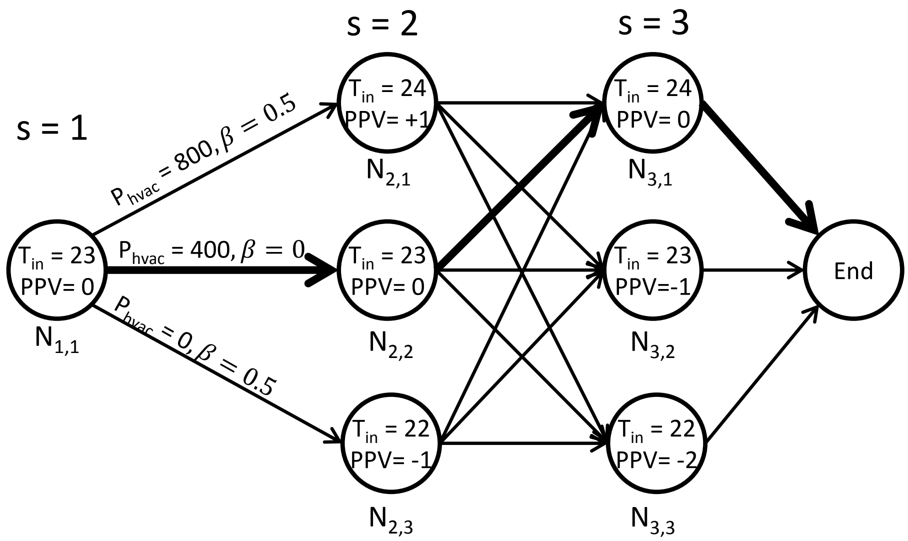

Figure 2: We use a graphical state model to find the optimal control sequence. Each state in the graph represents a potential control outcome at a ten-minute time step s. T_{in} is the predicted indoor temperature and PPV is the Predicted Personal Vote given such an indoor temperature. Note that we are making the implicit simplifying assumption that future comfort values are functions of temperature alone, i.e., that there is no change in worker clothing levels or in the environment variables. For each edge, P_{hvac} is the state transition energy cost, and β is the comfort penalty of the destination state. The optimal control sequence is the shortest path from the leftmost state to the rightmost state. The graph is recomputed at the end of each time step.

Note that the heater we used only supports on/off control, while our model in §3.3 assumes a heater with continues input power. To simulate a heater with variable heating power, we use pulse width modulation (PWM). Let the input power of the heater be P_{max} Watt and we defined a control cycle of C seconds. To simulate the input power of P_{hvac} Watt, the heater is turned on for $\frac{CP_{hvac}}{P_{max}}$ seconds in the control cycle and off for the rest of time in that cycle.

4.1 Temperature Control Schemes

We now describe the five different temperature control schemes that we implemented in this workspace.

- **Fixed Setpoint:** This control scheme has a fixed temperature setpoint of 25°C. If the measured temperature is lower than 25°C, the heater heats the room until measurements indicate that it has reached the setpoint.
- **Scheduled Setpoint:** This emulates the behaviour of a "Smart Thermostat": the controller maintains the room temperature at 25°C from 8 AM to 6 PM.
- **Reactive Temperature:** This control scheme starts to heat the room when occupancy is detected, and maintains a setpoint of 25°C only when the worker is present. To improve the robustness of the system, heating commences only after 5 minutes of continuous occupancy and stops when the workspace is vacant for 5 minutes. This reduces sensitivity to transient occupancy.
- **Reactive PPV (SPOT):** Instead of maintaining a constant temperature as in reactive temperature control, reactive PPV control maintains personal comfort (PPV) at the category B thermal comfort environment [3] where $ppv \in [-0.5, 0.5]$.
- **Optimal:** This scheme finds the best heating control sequence using LBMPC. We maximize worker comfort by setting λ in Eq. 14 to 10000 and ϵ to 0.5.

We attempted to run each control scheme for at least two days; the actual number of days for each scheme is shown in Table 1. For reactive temperature control, we obtained two days of data but later discovered that one day's data was not valid because one sensor had stopped working. Therefore the results for this scheme are not reliable. On the other hand, we have four and five days data respectively for reactive PPV and optimal control schemes, so the comparison of their relative performance is more reliable. Recall that all experiments were done in a building where the existing heating system maintained temperature at very nearly 23°C independent of the external temperature, so the results from the different days are comparable. Nevertheless, we stress that these results are far from statistically valid, and therefore must be viewed as suggestive, rather than definitive.

4.2 Evaluation Metrics

We propose three performance metrics. The **Average Daily Energy Consumption** is the total average energy consumed over a day. The **Average Absolute PPV** is the average absolute PPV value conditional on occupancy (because personal comfort only matters if the workspace is occupied). If $ppv(t)$ denotes the PPV value for time slot t

Control Scheme	Number of Days
Fixed	2
Scheduled	2
Reactive Temperature	1
Reactive PPV	4
Optimal	5

Table 1: Number of days tested for each control scheme.

during the day, the average absolute PPV is:

$$\frac{\sum_{t=1}^{T} ppv(t)m(t)}{\sum_{t=1}^{T} m(t)} \tag{17}$$

where $m(t)$ is 1 when the workspace is occupied and 0 otherwise.

Consider a control scheme A that always maintains the PPV at -0.5 and an alternative control scheme B that maintains PPV at 0 for half of the time and -1 for the other half. Both schemes have the same average absolute PPV. However, a worker will feel much more comfortable under scheme A because a typical worker is comfortable in the PPV range of $[-0.5, 0.5]$. Thus, we define the **Average Discomfort** to quantify how uncomfortable a worker feels over a day. Specifically, we define the discomfort at timeslot t as:

$$d(t) = max(|ppv(t)| - 0.5, 0) \tag{18}$$

In other words, if the PPV at timeslot t is in $[-0.5, 0.5]$, the discomfort d_t is 0, otherwise, the discomfort is $|ppv_t| - 0.5$. We then calculate the average discomfort conditional to occupancy as in Eq. 17.

4.3 Evaluation Results

4.4 LBMPC

We first evaluated the accuracy of LBMPC controller's thermal model by choosing a variety of setpoints between 21°C and 27°C and determining whether the temperature predicted the model matched the actual temperature. Over a period of two days, we found that there was a good match between these values, with a root-mean-square error of only 0.22°C. This shows that the LBMPC control is feasible for the office room under study.

4.4.1 Comparision between control schemes

We compare the five temperature control schemes based on the three metrics discussed above. Figure 3 and 4 show the results.

The fixed temperature setpoint control consumes 21.08 kWh of electricity daily and has an average absolute PPV and average discomfort of 0.42 and 0.07 respectively. Because the fixed temperature setpoint control always keeps the room temperature at 25°C, it guarantees the highest level of user comfort but consumes the most of energy as shown in the figures.

By using a temperature control schedule, we can save about half of the energy consumed, because the heater is off at night (i.e. from 5:30pm to 8:30am). Thus, we see that scheduled setpoint control consumes only 11.6 kWh of energy. However, this comes at a cost: if the worker is present in the workspace at an unscheduled time, or when the room is not sufficiently heated, the comfort target is not met. Thus, this scheme's average absolute PPV and average discomfort are higher at 0.71 and 0.23 respectively. One could argue that the comfort target could be met by choosing a more conservative schedule. This is, of course, true, but the choice of a more conservative schedule would increase the energy cost. Further, without knowledge of the room's thermal characteristics, it is difficult to determine precisely how early to start heating the room before the worker's arrival time.

We can further reduce energy consumption using *reactive* temperature control. The worker regularly leaves the workspace during working hours for lunch and meetings. Reactive control turns the heater off during these occupancy gaps to save energy. Over a day, the reactive control system consumes only 6.30 kWh of electricity (a further reduction of about 40%), but, because the workspace is cooler than desirable at the start of the day, and because the reactive temperature control scheme controls temperature, not worker comfort, its average absolute PPV and average discomfort are 0.86 and 0.36 respectively.

Reactive PPV control, as in SPOT, by maintaining a constant PPV level rather than a temperature setpoint, can improve user comfort. We see that it increases the user comfort with nearly no change to energy consumption. The daily energy consumption of reactive PPV control is 5.04 kWh (a slight decrease, actually), but its average absolute PPV and average discomfort are 0.53 and 0.20 respectively, which are significant improvements.

Recall that we set the tuning parameter for optimal control to greatly weight user comfort. Thus, the scheme maximizes user comfort without great regard to energy cost. Nevertheless, we find that optimal control consumes about 7.62 kWh of electricity (a slight increase over reactive control) but its average absolute PPV and average discomfort are 0.47 and 0.02, which are comparable to the fixed scheme.

In summary, we find that optimal control provides nearly the best comfort of all the schemes, but does so while reducing energy cost by a factor of more than three compared to fixed setpoint control, and a factor of more than two compared to scheduled setpoint control. It uses roughly the same energy as a reactive control scheme, but reduces discomfort by an order of magnitude (0.2 to 0.02). Thus, it has the best tradeoff between energy consumption and user comfort.

We now present a more detailed comparison between the reactive PPV and optimal scheme. Figure 5 and 6 show a typical example of the behaviour of the two schemes. Each figure shows the PPV value over the course of the day, with the bars on the X axis representing occupancy. For reactive control, the room starts to heat at around 8:40 AM when the worker arrives. It requires more than half an hour to reach the target PPV value. In Figure 2, the room starts to pre-heat at around 7:30 AM and when the occupant arrives at 9:00 AM, the room temperature is already in the comfort zone. This demonstrates how the optimal scheme improves user comfort over the reactive scheme by forecasting workspace occupancy. Note also that the optimal scheme, by frequently re-evaluating the optimal heater control sequence, is able to adjust automatically to unpredicted changes in workspace occupancy.

5. RELATED WORK AND DISCUSSION

There has been extensive work both on HVAC modeling [14] as well as optimal HVAC control in both residen-

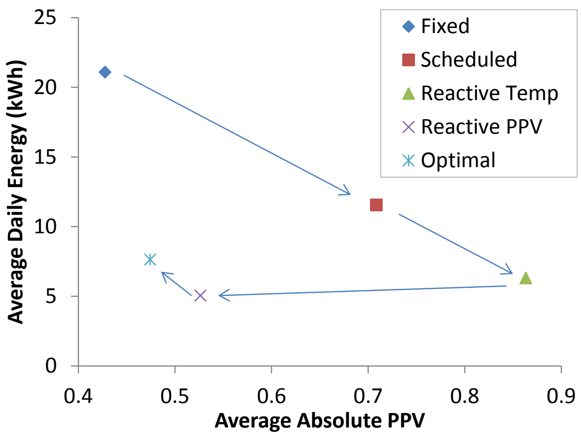

Figure 3: Average absolute PPV versus daily energy consumption for five control schemes. The fixed (always-on) scheme has the highest energy cost but the best performance. The scheduled control scheme reduces the energy cost but increases discomfort. The reactive scheme based on temperature has a low energy cost, but causes the most discomfort. The reactive scheme based on PPV has nearly the same energy cost as the one based on temperature, but lowers the energy cost. Finally, the optimal scheme achieves nearly the same comfort level as the always-on scheme, but at a much lower energy cost.

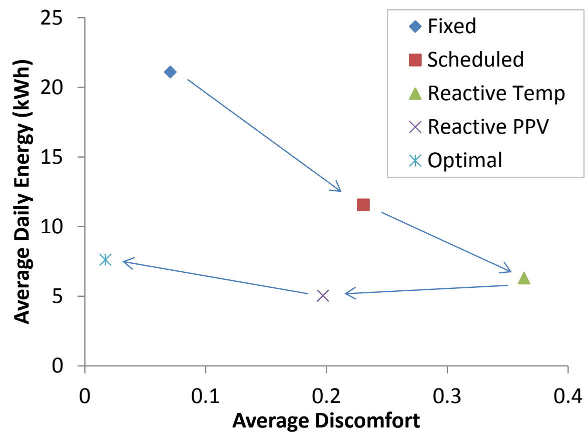

Figure 4: Average discomfort versus daily energy consumption for five different control schemes. The trends in this metric are nearly identical to those in Figure 3.

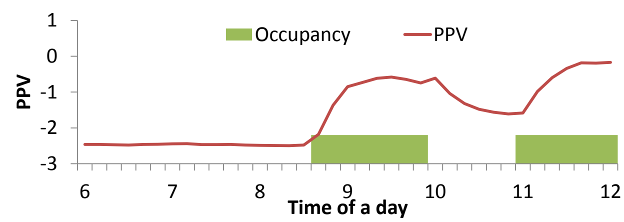

Figure 5: PPV vs. time of day for the reactive PPV control scheme. Note that this control starts to heat the room when occupancy is detected.

tial and commercial buildings, as surveyed in Reference [17]. There has also been some prior work on *comfort index regulators* that control HVAC systems so as to maintain a specified level of thermal comfort for occupants (for example, see References [4,12] and the work cited therein). Unlike this prior work, which typically carefully model a central HVAC plant and use optimization as the solution approach, our focus on *personal* comfort allows us to use a particularly simple (i.e., single zone) physical model and a solution approach based on optimal control.

The work that is most closely related to ours is by Gomez-Otero et al [9], where the estimated PMV in a room is used to control the setpoint and airflow of individual HVAC units. Our work differs from their in three ways. First, they control room cooling, rather than room heating. Second, we use a Kinect for clothing level computation as well as occupancy ground truth. In contrast, they use heuristics for clothing level estimation and a smartphone for worker localization. Finally, we provide exact energy costs, whereas in their work, they are only able to estimate energy efficiency from the number of hours of operation of the HVAC units. An additional problem with their work is that they do not describe their algorithms or control scheme in any detail, making it impossible to compare their work with ours.

Aswani et al [4] have used an LBMPC approach for HVAC control, but our use of the Microsoft Kinect sensor to establish ground truth about workspace occupancy, which is novel, allows us to eliminate a significant source of control error.

PreHeat [16] also attempts to improve user comfort by starting heat a room in anticipation of future occupancy. We differ in our use of LBMPC and optimal control to accurately calculate the optimal time to start heating. Moreover, we use the PPV model to accurately model user thermal comfort, which indicates the upper bound for temperature setback.

Our experiments indicate that the quality of the optimal control schedule is limited by the accuracy of occupancy prediction. The current accuracy of occupancy prediction at a half-hour granularity is about 80% during the day time (6 AM to 6 PM). By increase the time granularity to 10 minutes, we can improve the accuracy of occupancy prediction to almost 90%. We find it difficult to further improve the accuracy unless we can obtain more side channel information. For example, Gupta et. al. [10] use a GPS mounted on the home owner's car to forecast the arrival time of the home owner. Similarly, Ardakanian et. al [5] use the sound and light level of a room to infer occupancy and use a Partially-Observable Markov Decision Process (POMDP) for building heating control. We anticipate incorporating these into our future work.

6. CONCLUSION

Building on our prior work on personal thermal control [8], we have described SPOT+, an LBMPC-based optimal control framework that find the optimal balance between the energy-comfort tradeoff. Using occupancy and temperature prediction, SPOT+ finds the best control schedule that min-

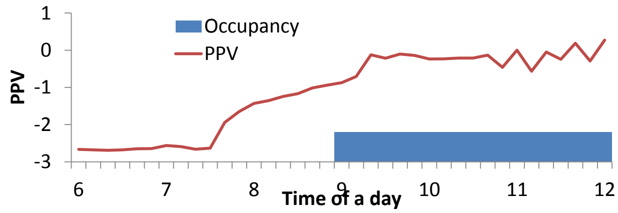

Figure 6: PPV vs. time of day for the optimal control scheme. Note that this control starts to heat the room *before* the estimated arrival time

imizes the energy consumption without affecting user thermal comfort. Preliminary experiments show that SPOT+ saves about 60% of energy and reduces user discomfort from 0.36 to 0.02 comparing with different baseline methods.

In future work, we plan to use our approach to provide personal thermal comfort in open-plan office spaces. We are also considering the tradeoff between using low-power heaters that heat for a long time but do not cause load peaks versus high-power heaters that can quickly heat a space, but can cause load peaks.

7. REFERENCES

[1] Buildings and Communities. *http://canmetenergy.nrcan.gc.ca/buildings-communities/2218*.

[2] Energy Efficiency Fact Sheet: Heating, ventilation & Air Conditioning. *http://www.originenergy.com.au/files/SMEfs_HeatingAirCon.pdf*.

[3] 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*.

[4] 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.

[5] O. Ardakanian and S. Keshav. Using Decision Making to Improve Energy Efficiency of Buildings. *ICAPS-10 POMDP Practitioners Workshop, May 2010*.

[6] R. J. de Dear and G. S. Brager. Thermal comfort in naturally ventilated buildings: revisions to ashrae standard 55. *Energy and buildings*, 34(6):549–561, 2002.

[7] P. O. Fanger. Thermal comfort. *Analysis and Applications in Environmental Engineering, Danish Technical Press, Copenhagen, Denmark, 1970*.

[8] P. X. Gao and S. Keshav. SPOT: A Smart Personalized Office Thermal Control System. In *Proc. ACM e-Energy '13*, 2013.

[9] C. Gomez-Otero, R. Martinez, and J. Caffarel. Climapp: A novel approach of an intelligent hvac control system. In *Information Systems and Technologies (CISTI), 2012 7th Iberian Conference on*, pages 1–6. IEEE, 2012.

[10] 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.

[11] International Energy Agency. Technology roadmap: Energy-efficient buildings: Heating and cooling equipment. 2011.

[12] J. Liang and R. Du. Design of intelligent comfort control system with human learning and minimum power control strategies. *Energy Conversion and Management*, 49(4):517 – 528, 2008.

[13] G. K. C. Mehdi Shahrestani, Runming Yao. Performance Characterisation of HVACR Systems for Building Energy Benchmark. *http://www.cibse.org/content/cibsesymposium2012/Poster026.pdf*.

[14] L. Perez-Lombard, J. Ortiz, and I. R. Maestre. The map of energy flow in HVAC systems. *Applied Energy*, 88(12):5020 – 5031, 2011.

[15] L. Perez-Lombard, J. Ortiz, and C. Pout. A Review on Buildings Energy Consumption Information. *Energy and Buildings*, 40(3):394 – 398, 2008.

[16] 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.

[17] S. Wang and Z. Ma. Supervisory and Optimal Control of Building HVAC Systems: A Review. *HVAC&R Research*, 14(1):3–32, 2008.

APPENDIX

Here are the two tables used for PMV computation from the ISO 7730 standard.

Activity	**Metabolic Rate**	
	W/m^2	*met*
Seated, relaxed	58	1.0
Sedentary activity	70	1.2
Standing, medium activity	93	1.6

Table 2: Metabolic Rates

Work Clothing	**Clothing Insulation** (I_{cl})	
	clo	$m^2 \cdot K/W$
Underpants, shirt, trousers, smock, socks, shoes	0.90	0.140
Underwear with short sleeves and legs, shirt, trousers, jacket, socks, shoes	1.00	0.155
Underwear with long legs and sleeves, thermojacket, socks, shoes	1.20	0.185
Underwear with short sleeves and legs, shirt, trousers, jacket, heavy quilted outer jacket and overalls, socks, shoes, cap, gloves	1.40	0.220
Underwear with short sleeves and legs, shirt, trousers, jacket, heavy quilted outer jacket and overalls, socks, shoes	2.00	0.310

Table 3: Thermal Insulation for different clothing levels