Directional Airflow for HVAC Systems

Milad Abedi

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Civil Engineering

Farrokh Jazizadeh Karimi, Chair
Bert Huang
Pratap Tokekar

December 11, 2018
Blacksburg, Virginia

Keywords: HVAC, Thermal Comfort, Reinforcement Learning, Building Energy Consumption

Directional Airflow for HVAC Systems

Milad Abedi

(ABSTRACT)

Directional airflow has been utilized to enable targeted air conditioning in cars and airplanes for many years, where the occupants could adjust the direction of flow. In the building sector however, HVAC systems are usually equipped with stationary diffusors that can only supply the air either in the form of diffusion or with fixed direction to the room in which they have been installed. In the present thesis, the possibility of adopting directional airflow in lieu of the conventional uniform diffusors has been investigated. The potential benefits of such a modification in control capabilities of the HVAC system in terms of improvements in the overall occupant thermal comfort and energy consumption of the HVAC system have been investigated via a simulation study and an experimental study. In the simulation study, an average of 59% per cycle reduction was achieved in the energy consumption. The reduction in the required duration of airflow (proportional to energy consumption) in the experimental study was 64% per cycle. The feasibility of autonomous control of the directional airflow, has been studied in a simulation experiment by utilizing the Reinforcement Learning algorithm which is an artificial intelligence approach that facilitates autonomous control in unknown environments. In order to demonstrate the feasibility of enabling the existing HVAC systems to control the direction of airflow, a device (called active diffusor) was designed and prototyped. The active diffusor successfully replaced the existing uniform diffusor and was able to effectively target the occupant positions by accurately directing the airflow jet to the desired positions.

Directional Airflow for HVAC Systems

Milad Abedi

(GENERAL AUDIENCE ABSTRACT)

The notion of adjustable direction of airflow has been used in the car industry and airplanes for decades, enabling the users to manually adjust the direction of airflow to their satisfaction. However, in the building the introduction of the incoming airflow to the environment of the room is achieved either by non-adjustable uniform diffusors, aiming to condition the air in the environment in a homogeneous manner. In the present thesis, the possibility of adopting directional airflow in place of the conventional uniform diffusors has been investigated. The potential benefits of such a modification in control capabilities of the HVAC system in terms of improvements in the overall occupant thermal comfort and energy consumption of the HVAC system have been investigated via a simulation study and an experimental study. In the simulation study, an average of 59% per cycle reduction was achieved in the energy consumption. The reduction in the required duration of airflow (proportional to energy consumption) in the experimental study was 64% per cycle on average. The feasibility of autonomous control of the directional airflow, has been studied in a simulation experiment by utilizing the Reinforcement Learning algorithm which is an artificial intelligence approach that facilitates autonomous control in unknown environments. In order to demonstrate the feasibility of enabling the existing HVAC systems to control the direction of airflow, a device (called active diffusor) was designed and prototyped. The active diffusor successfully replaced the existing uniform diffusor and was able to effectively target the occupant positions by accurately directing the airflow jet to the desired positions.

Dedication

To my parent for their unwavering love and support

Acknowledgments

I would like to thank my advisor Dr. Jazizadeh for all of his support and guidance during my research. I would also like to thank my advisory committee Dr. Huang, and Dr. Tokekar for their kind support and guidance.
This material is based upon work supported by the Institute for Critical Technology and Applied Science (ICTAS) at Virginia Tech. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ICTAS.

Attribution

Research Study 1: In this study, Dr. Francine Battaglia and Dr. Farrokh Jazizadeh developed the CFD Model of the test room and performed the CFD simulation. Milad Abedi analyzed the output data of the CFD simulation and drafted the article. All the authors reviewed and made modifications to initial drafts. Dr. Farrokh Jazizadeh presented the article in the "25th International Workshop on Intelligent Computing in Engineering".

Research Study 2: In this study, Dr. Francine Battaglia and Dr. Farrokh Jazizadeh developed the CFD Model of the test room and performed the CFD simulation. Milad Abedi analyzed the output data of the CFD model, developed and evaluated the Reinforcement Learning framework and drafted the article. Dr. Francine Battaglia and Dr. Farrokh Jazizadeh developed the CFD Model of the test room and performed the CFD simulation. All the authors reviewed and made changes to initial drafts.

Research Study 3: In this study Milad Abedi, under the supervision of Dr. Farrokh Jazizadeh, designed and built the active diffusor. Under the supervision of Dr. Farrokh Jazizadeh, Milad Abedi developed the experimental setup, performed the experiments, analyzed the results and drafted the article. Dr. Jazizadeh reviewed and modified the initial draft of the article.

Contents

1 Introduction

Numerous studies have endeavored to improve the performance of Heating, Ventilation, and Air-Conditioning (HVAC) systems in terms of occupant thermal comfort and energy expenditure. In the US, buildings are responsible for consumption of 39% [1] of the total energy. HVAC system are responsible for 44% [2] of the energy consumed in commercial buildings and 48% [3] of the energy consumption in residential buildings, resulting in high incentives for any effort towards higher energy efficiency of HVAC systems.

In the traditional approach to HVAC control, the control process is designed around the notion of the HVAC systems' responding to the temperature set-point requirements of the controlling thermostat. In so doing, a number of unwarranted assumptions have been made with regards to occupant comfort tendencies and thermal zone properties. Firstly, the ambient temperature at the location of the controlling thermostat is not necessarily the same as the ambient temperature at the vicinity of the occupant.

In an attempt to render the existing HVAC systems more responsive to the thermal comfort needs of occupants while avoiding intrusive modifications of the inner operations and the control logic of the existing HVAC systems, researchers [4–6] have proposed control frameworks where the temperature set-point of the controlling thermostat would be driven in accordance with the thermal comfort preferences of the occupants. Notwithstanding the relative success of this strategy in improving the conventional control approach, further improvements are possible. For instance, it is a common occurrence that one thermostat is used to control the HVAC system dedicated to a zone consisting of multiple rooms. In this scenario, in response to thermostat temperature set-point requirements, the HVAC system will direct the air to all rooms in the zone regardless of their state of occupancy, resulting in unnecessary energy loss due to directing the conditioned air to rooms where there is no need for it.

An intuitive solution (called zoning) to the aforementioned short-coming is to enable room-level binary control in which, the HVAC system would be able to open and close the air-vent to each room in accordance with requirements of thermostats installed in each room, thereby resulting in a further level of control granularity. While the level of control achieved by the zoning solution allows for energy savings, the zoning solution still views each room as a uniform entity and does not take into account the variations in temperature distribution through the room.

In this thesis, in an effort to account for the non-uniformity of the thermal properties in the room space, the idea of directional airflow has been investigated. While in the conventional HVAC systems the conditioned air is diffused uniformly through the room environment, in this thesis the notion of directional airflow has been investigated as an improvement upon the traditional uniform diffusion approach. By utilizing directional airflow, a local area of comfort can be created in the vicinity of the occupant, thereby avoiding the need to condition the air in the entirety of the space, resulting in savings in energy expenditure.

This thesis consists of three studies as follows. In the first research study, a Computational Fluid Dynamic (CFD) model of a typical office room was augmented by machine learning algorithms to study the potential benefits of directional airflow on HVAC energy consumption and overall occupant thermal comfort in a simulation setting. In the second research study, the feasibility of autonomous control of the direction of airflow at the point of the ceiling air vent, subject to privacy constraints was investigated. To do so, outputs of a CFD model of a typical office room were used to build a simplified model of the thermal environment of the room, in which the effect of the directional airflow could be simulated. The Reinforcement Learning (RL) algorithm was utilized as the control algorithm, operating in the absence of knowledge about the location of the occupant to obviate the need for utilization of intrusive indoor localization technologies. In the third research study, a robotic device (called the active diffusor) was designed and built that would allow the HVAC system to control the direction of airflow. The active diffusor has been designed such that it could readily replace the existing uniform diffusors. In the third paper, the feasibility of the active diffusor has been studied in terms of its capability in effectively directing the airflow jet towards the direction of interest and the energy saving potential of utilization of such a device.

2 Research Study 1: Smart HVAC Systems — Adjustable Airflow Direction

Authors: Milad Abedi[1], Farrokh Jazizadeh[1], Bert Huang[2], Francine Battaglia[3]

[1]Via Department of Civil and Environmental Eng., Virginia Tech, Blacksburg, VA, 24061, USA
[2]Department of Computer Science, Virginia Tech, Blacksburg, VA, 24061, USA
[3]Department of Mechanical & Aerospace Engineering, University at Buffalo, NY, 14260,USA

Key Words: Airflow Direction, HVAC, Thermal Comfort, Computational Fluid Dynamics, Statistical Learning

2.1 Introduction

Thermal comfort, defined as "that condition of mind which expresses satisfaction with the thermal environment and is assessed by subjective evaluation" [7], has been the subject of many research efforts during the last decades. While the comfort itself is a valid objective for research, it has also been established that increase in thermal comfort often leads to decrease in energy consumption. This is particularly encouraging due to the considerable contribution of HVAC systems to the total energy consumption. In the US, buildings are responsible for 40% of the total energy consumption [1], and 48% of the energy consumed in residential buildings is used by HVAC systems [3]. In commercial buildings, HVAC systems are responsible for 44% of the total energy consumption. [2] Current thermal comfort models such as the Predicted Mean Vote model [7] depend on a set of environmental and human parameters such as air temperature and metabolic rate. However, in real-world applications, lack of information about the specific occupant and the environment often leads to conservative assumptions. These assumptions in return lead to sub-optimal HVAC performance. Several research efforts [4, 8–18] have focused on personalized HVAC control systems so as to adjust the HVAC system according to the occupant-specific thermal comfort tendencies to avoid making overly general assumptions. The other aspect of the problem, namely the estimation of relevant environmental parameters also presents a challenge especially considering the infeasibility of installing large number of sensors to better understand the environment. With the low number of sensing points, the condition of the environment is assumed to be one of relative uniformity thereby failing to take into account the effects of a non-uniform environment on the occupant's thermal comfort. Lack of sufficient awareness of the environment and the consequent simplistic assumptions also affect the actions taken by HVAC system to improve the occupant's thermal comfort. Aside from the lack of information about the environment, the inflexibility of HVAC system in terms of possible actions also limits the HVAC system's ability to produce an effective, environment-aware response. For instance, in a room with several diffusors, a change in one diffusor's airflow rate might be a sufficient action to ensure the thermal comfort of the occupant but the HVAC system's inability to control every single diffusor would limit the control to a change in airflow rate of all the diffusors which can be an ineffective action in terms of energy consumption. In our previous studies [4, 12–17] we have established methods of evaluating thermal comfort of occupants that could address the issue of lack of sufficient environment awareness in terms of its effect on evaluating the occupant's thermal comfort state. These methods also account for the natural differences in the occupant's thermal comfort parameters, thereby avoiding conservative assumptions regarding the parameters in question, which can in return lead to higher thermal satisfaction as well as lower energy consumption. In so doing, the occupant could be used as proxy for sensing points. One of the potential merits of using the occupant as a proxy for sensing points is that the HVAC system might be able to take more specific actions to satisfy the occupant with lower energy consumption. Returning to the previous example of a room with multiple diffusors, having a real-time understanding of the occupant's thermal comfort state can help the HVAC control algorithm take more effective actions. To address the problem of lack of environmental information, we can use Artificial Intelligence (AI) learning algorithms that explore the environment and learn as they interact with it thereby alleviating the need for prior knowledge of the environment. One of such methods is Reinforcement Learning, where an agent learns how to behave in an unknown environment by interacting with it. While the current state of the literature on flexible HVAC systems in terms of control within a thermal zone is very limited, a fair amount of research exists on a building level flexibility of HVAC systems. Y. Tachwali et al. [19] designed a multi-zone HVAC controller equipped with a wireless sensor network that was capable of directing the airflow only to the occupied zones. In other studies, S.R West et al. [20] and F. Jazizadeh et al. [4] enabled a zone-level temperature set-point control for an HVAC system with the objective of further reducing running costs and CO2 emissions while maintaining occupant's thermal comfort. Through simulation of an entire floor with several rooms and offices and by using the occupancy data gathered from monitoring the real-world counterpart of the simulated model, Y. Agarwal et al. [21] studied the benefits of a zone level temperature set-point control based on the occupancy state of each zone. In a subsequent paper, Y. Agarwal et al. [22] explored the idea of using a network of occupancy detection sensors to enable a zone-level airflow control (on or off). Benefits of control over the components of the HVAC system has also been explored. N. Nassif et al. [23], in addition to the 70 zone-temperature set-points also controlled the inner components of the HVAC system such as supply duct static pressure, supply air temperature and chilled water supply temperature. In their study, M. Feldmeier et al. [9] extended the level of control beyond the HVAC system, which was able to control zone-level temperature set-points, and enabled the control algorithm to open and close windows as well. While the smart control of the airflow direction by the HVAC system has not been properly studied, M. Fountain et al. [24] explored the benefits of allowing the users to control the direction of flow at the point of diffusion and also by using a desk fan. They concluded that enabling control over the air-movement can lead to higher thermal satisfaction. As a first step in the way to designing and building more effective HVAC system, in this paper we have sought to answer the question of whether increasing the flexibility of the HVAC system in terms of control over the direction of airflow could lead to higher overall occupant thermal satisfaction with lower energy consumption. To do that a Computational Fluid Dynamic (CFD) model of a room was developed and used to study the effects of direction of airflow on the performance of the HVAC system. The performance of such a system was evaluated in term of energy consumption. The results of the analysis were used to gain insight into the promises of such a system and the challenges that the envisioned HVAC system would have to overcome.

2.2 Methodology

As noted, the objective of this paper is to investigate the effect of adjusting the direction of airflow on the energy consumption of the HVAC system while taking into consideration the thermal comfort state of the occupant. This objective has been pursued through a coupled simulation consisting of physics-based modeling and machine learning-based prediction algorithms. The simulation represents an occupant sitting in a room equipped with an HVAC system with adjustable airflow direction.

2.2.1 Environment

A CFD model in the commercial software ANSYS Fluent [25] was developed to serve as the simulation environment. The model represents a $6m \times 4m \times 2.5m$ room with a single diffusor on one side and an exhaust on the other (Figure 1 (a)). Given the high computational cost of implementing the CFD model, the direction of the airflow was set to have a constant 30° angle with the ceiling so that the space of possible directions would become smaller and easier to explore. Equations for conservation of mass, momentum and energy were considered for an incompressible Newtonian fluid, and the Boussinesq model and the $k - \epsilon$ turbulence model were incorporated. The equations were solved by using the segregated Pressure-Based Navier-Stokes (PBNS) numerical solver and the Semi-Implicit Method for Pressure-Linked Equation (SIMPLE) algorithm was employed to solve the pressure-velocity coupling. Gradient and pressure spatial discretization was performed through the least squares cell based (LSCB) and the pressure staggering option (PRESTO!) approaches, respectively. For time derivative discretization, a first-order implicit method with a time step of 0.1 seconds was used. Discretization of energy, momentum, dissipation rate, turbulent kinetic energy and discrete ordinate was performed using a second-order upwind. Previous studies by Park et al. [26] and Wang et al. [27] have validated the use of these equations and models for natural and forced convection in buildings. Through a grid resolution study the uniform grid spacing of 14.2 cm was chosen. A detailed validation of the grid cell size was performed in the studies by Park et al. [26, 28] and guided the cell size used here. The room is exposed to the warmer outside environment with a constant temperature of 30 °C (corresponding to a hot summer day) through a wall on one side of the room (Figure 1). The solar intensity is constant at 1100 W/m^2, which represents a southern-facing window at a latitude of 37°N. The window shades are assumed to reflect solar radiation, but heat transfer is allowed through the windows. The stone façade of the exterior office walls is made of dolomite (thermal conductivity is 1.5 W/mK) and conduction is modeled through the exterior walls. The diffusor has been modeled using conventional 4-way louvre-bladed diffusors where the adjustable nature of the diffusor is defined as boundary conditions using velocity, flow area opening, and corresponding blade angles in each of the directions. The incoming airflow from the diffusor has a temperature of 13 °C and can be simulated at any direction and for any duration at the rate of 200 cfm (0.0944 m^3/s). In order for the CFD model to be representative of real-world scenarios the CFD model was started from an initial condition and it was allowed to run for 180 seconds without any diffusion. During this period, the room exchanged heat with the outside environment and thus a temperature gradient was created in the room. After the initial 180 seconds, the diffusor started to work and we continued to run the model with the diffusor turned on for an additional 420 seconds, while recording the results at 10 second time steps.

2.2.2 Occupant

For simplicity, we have assumed that the simulated room will only host one occupant. The occupant is assumed to be seated in a chair, the top of the head of the occupant being at a height of 1.2 meters from the floor.

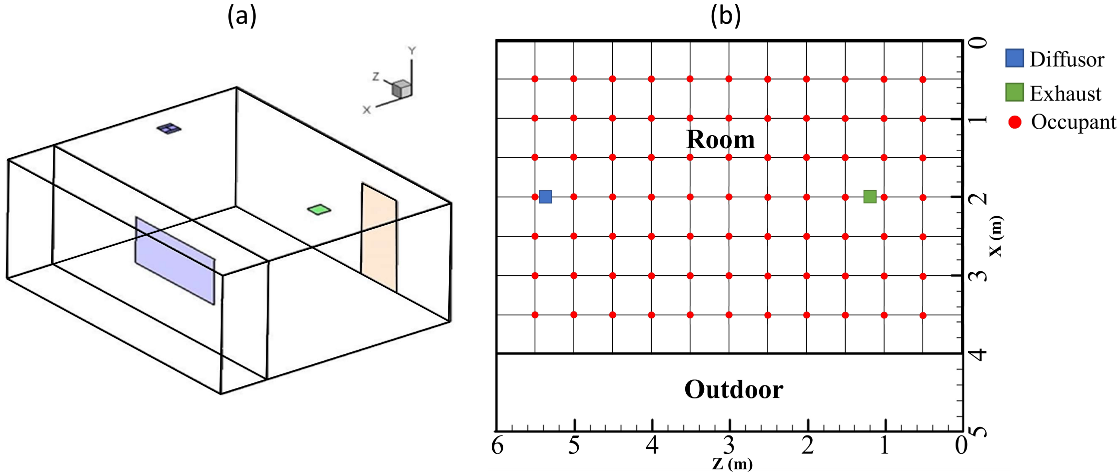

Figure 1: (a) 3-D sketch of the simulated room. (b) The plan of the room in the CFD simulation. The considered occupant locations have been identified by red dots.

The occupant is assumed to be stationary. The computations have been carried out for 77 different locations for the occupant, corresponding to a mesh of 11 by 7, with 50cm steps (Figure 1 (b)). This configuration corresponds to assuming a uniform probability distribution for the coordinates of the occupant in the room. To evaluate the thermal comfort state of a simulated occupant at each step, a scaled version of the Proportion Comfortable model presented by Daum et al. [18] (Figure 2 (a)) was used. This model presents a relationship between the surrounding temperature and the thermal comfort index of the occupants. For this study, the aggregate thermal comfort graph was normalized so that its maximum equals 1 (Figure 2 (b)).

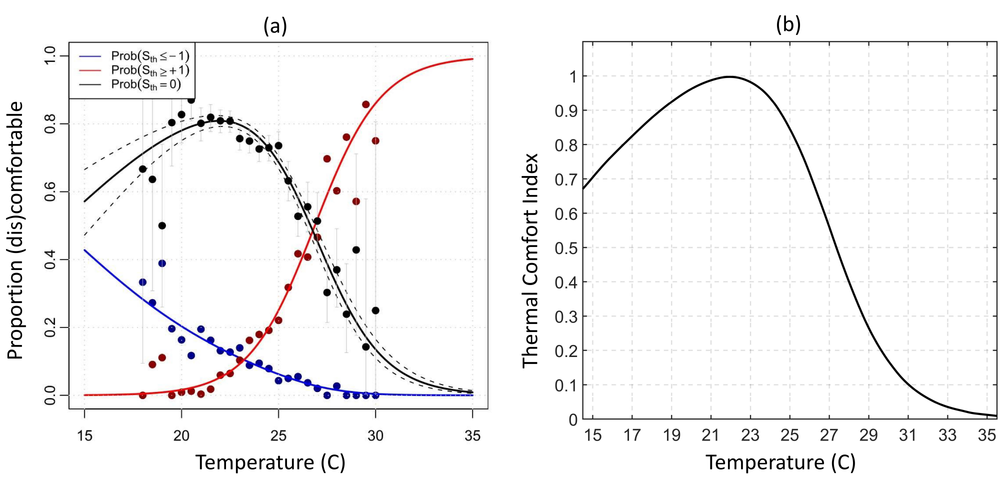

Figure 2: Thermal comfort index versus temperature. (a) The original model taken from [18]. (b) The scaled model.

By using this model, the thermal comfort index can be calculated based on the temperature in the vicinity of the occupant. At each time step of the CFD simulation, the temperature surrounding the occupant is calculated by averaging the temperature at three heights corresponding to the head, abdomen and ankles of the occupant. The average temperature is then fed into the aforementioned model, which in return outputs the thermal comfort index of the occupant.

2.2.3 Using Machine Learning Algorithms to Expand the Results of the CFD Simulation

Given the prohibitive computational cost of running the CFD model for every airflow direction of interest, simulations were carried out only for a total number of 6 airflow directions as illustrated in Figure 3. A number of machine learning algorithms (Regression Models) were used to expand the results of the simulated 6 airflow directions to non-simulated directions of interest. Based on an empirical validation process, the Gaussian Regression Process (GPR) model was chosen. Using the CFD simulation results, a GPR model was trained that would allow us to evaluate the changes in the occupant's thermal comfort index in cases where the direction of airflow was not among the 6 simulated directions.

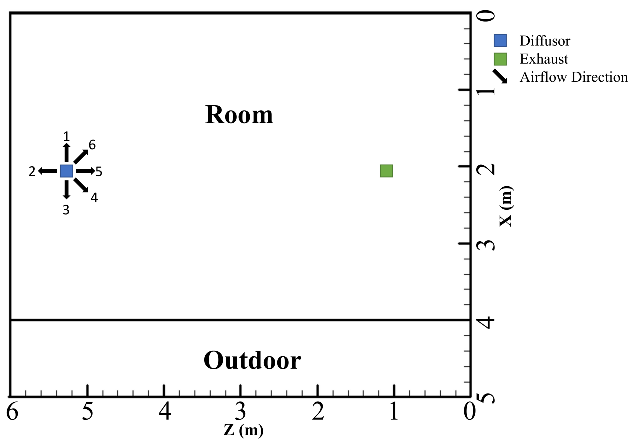

Figure 3: The simulated airflow directions. All directions have a 30° angle with the ceiling.

2.2.3.1 Model Selection and Validation Given that the intention behind using a regression model was to estimate the improvements in thermal comfort index in non-simulated directions, we have tailored our validation approach to correspond to this objective. To validate the performance of any regression model we completely eliminated the data belonging to a single chosen direction from the training set and used it as testing data instead. For instance, we trained the model based on the simulation results of directions 1, 2, 3, 5, 6 (leaving out direction 4). We then used the regression model to predict the results of the simulation for direction 4, and compared the regression model's predictions to the CFD simulation results. This process was performed in each direction for all the 77 occupant locations as shown in Figure 1 (b). Table 1 represents the average and minimum values for the coefficient of determination (R^2) for the testing data in each of the tests in the case of the GPR model. For instance, the first row of Table 1 represents the validation results for the case where the CFD simulation data for direction 1 were eliminated from the GPR model's training set. The R^2 values were calculated based on the testing data (belonging to direction 1) and predictions of the GPR model that was trained on the data for remaining directions.

Selection of the regression model to generalize the results of the CFD analysis was done through an empirical validation process. The metric that was used for model selection is the average of all the R^2 values. For instance, the value of this metric for the GPR model is the average of all the values in the second column of Table 1. A number of well-known regression methods were tested, and based on the overall validation results (Table 2) the GPR model was selected. The Boosted Decision Tree Regression model, despite its higher R^2 score was not chosen for reasons that have been explained at the end of this section. In Table 1, note the lower value of the average and minimum R^2 in the case where the data belonging to direction 2 was eliminated from the training data. This is not surprising, because once we eliminated the data belonging to direction 2 there was no airflow direction in the training data that had a component in the same direction as direction 2 (all the other directions have an angle of 90° or more with direction 2). The results in Table 1 indicate that GPR is an accurate and reliable approach for our needs.

Table 1: Validation Results for the GPR Model

Eliminated Direction	**Average R^2**	**Minimum R^2**
1	0.97	0.93
2	0.60	0.47
3	0.91	0.87
4	0.97	0.92
5	0.93	0.87
6	0.98	0.91

Table 2: Model selection based on average overall R^2

Regression Method	**Average Overall R^2**
Boosted Decision Tree Regression	0.98
Gaussian Process Regression	0.89
K Nearest Neighbors Regression (K=5)	0.79
Multi-layer Perceptron Regression	0.58
Support Vector Regression	0.46
Linear Regression	0.45

$$f(x) = X^T W, \quad y = f(x) + \epsilon \tag{1}$$

Where w represents the model parameter vector and ε is the noise which is assumed to have a normal distribution as follows.

$$\epsilon \sim N(0, \sigma_n^2) \tag{2}$$

Assuming that the training data are independent, the likelihood is defined as:

$$p(y \mid X, W) = \prod_{i=1}^{n} p(y_i \mid x_i, W) = N(X^T W, \sigma_n^2 I) \tag{3}$$

Assuming a zero mean Gaussian prior probability with a covariance matrix of Σ_p we have:

$$W \approx N(0, \Sigma_p) \tag{4}$$

$$p(W \mid y, X) = \frac{p(y \mid X, W)P(W)}{P(y \mid X)} \tag{5}$$

After some mathematical manipulation, we would arrive at the following conclusion.

$$p(W \mid X, y) \sim N(\overline{W}, A^{-1}) \tag{6}$$

Where:

$$A = \frac{1}{\sigma_n^2} X X^T + \Sigma_p^{-1} \tag{7}$$

$$\overline{W} = \sigma_n^{-2} A^{-1} X y \tag{8}$$

To make predictions for a test case we would average over all possible parameter values w, weighted by their posterior probabilities. Thus the predictive distribution for $f_* \triangleq f(X_*)$ at X_* is given by:

$$p(f_* \mid X_*, X, y) = \int p(f_* \mid X_*, W)\ p(W \mid X, y) dW = N\Big(\frac{1}{\sigma^2} X_*^T A^{-1} X y, X_*^T A^{-1} X_*\Big) \tag{9}$$

The strength of the GPR models stems partly from their capability in utilizing the notion of kernels (for further readings see ref. [29]).

2.2.3.2 GPR Model Overview Inputs to the GPR model are the airflow direction, the duration of airflow, and the initial thermal comfort index of the occupant (Figure 4). The first two parameters, namely the direction and duration of airflow define the actions that the proposed HVAC system can take. The initial comfort index of the occupant encapsulates a number of human and environmental variables. The output of the GPR model is the amount of improvement in the thermal comfort index of the occupant as a result of the action that was taken by the HVAC system (Figure 4).

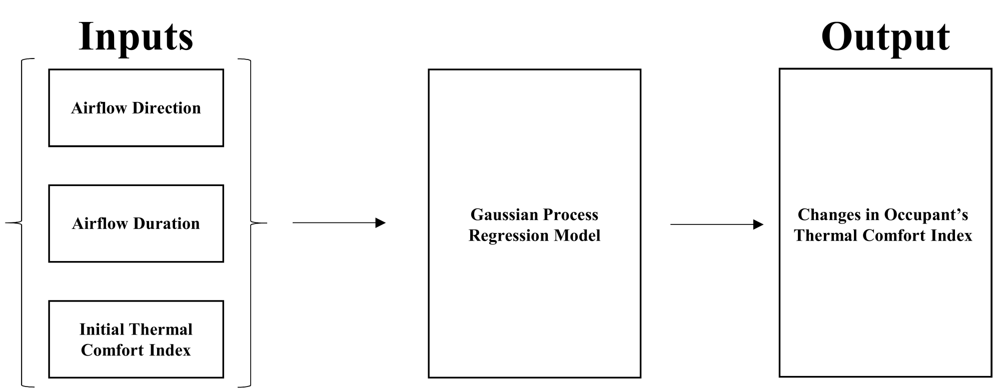

Figure 4: Overview of the GPR Model.

2.2.3.3 Evaluation of the Boosted Decision Tree Regression (BDTR) Model

While the higher value of the average R^2 metric in the first glance suggests a superior performance for the BDTR model, upon closer inspection, we came to the conclusion that the BDTR model might not be representative of the CFD simulation results. Our analysis indicated that the BDTR model might have over-fitted the model to duration of flow (one of the inputs to the model) and as a result may have not been able to represent the effect of airflow direction on improvements in the occupant's thermal comfort index. To show that the BDTR model does not capture the relationship between the direction of airflow and the model output, we trained the BDTR model by using the data belonging only to direction 2 and then used the trained model to predict the improvements in thermal comfort index for all the other directions. The average testing R^2 for this analysis was 0.98 which indicates that the results of the CFD simulation for airflow direction 2 was enough to predict the results for all the other direction, meaning that the airflow direction is inconsequential when it comes to improvements in the thermal comfort index of the occupant. If this were true, then the first hypothesis of this paper, namely that the direction of airflow has a significant effect on the thermal comfort index of the occupant would have been proven false. While R^2 is a well-known metric for validation or rejection of regression models, it is a single number and it does not provide enough insight into the underlying mechanism of the model. In order to gain further insight into the results of the BDTR model, we introduced the following error definition for any given data point in our data set.

$$Absolute\ Error_i = |PV_i - TV_i| \tag{10}$$

Where PV is the predicted value given by the regression model and TV is the true value obtained from the CFD simulation. To investigate the relationship between the BDTR model's performance and the Duration of Airflow parameter, a BDTR model and a GPR model were trained on the data belonging only to direction 2 and then they were used to predict the improvement in thermal comfort index for an occupant sitting at (X=3.5, Z=5.5) for all the other simulated airflow directions. Figure 5 (a) represents the absolute error of the BDTR model's predictions and its relationship with the duration of airflow.

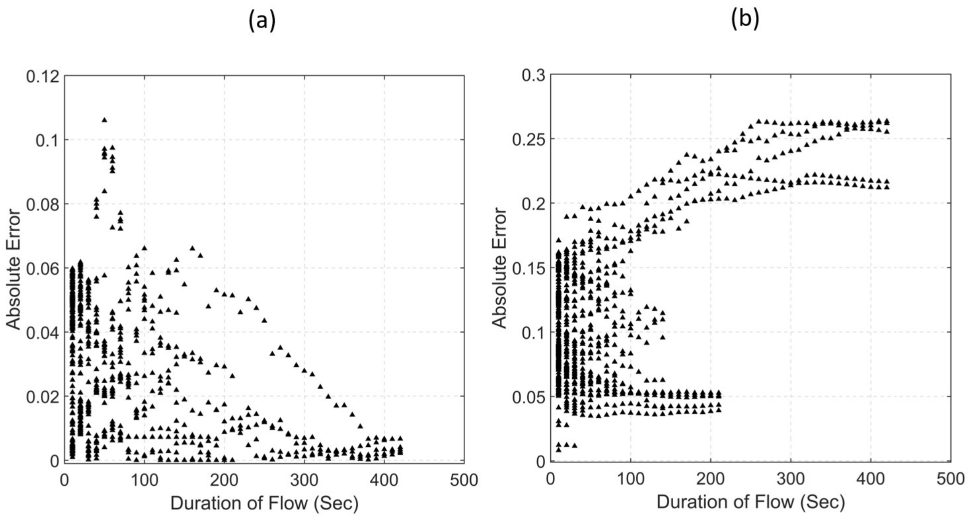

Figure 5: Absolute Prediction Error of regression models trained only on the data belonging to airflow direction 2 (a) BDTR model (b) GPR model.

As shown in Figure 5 (a), in the case of BDTR the data points that have a longer duration of flow have been predicted more accurately than those that have a lower duration of flow. This is not surprising, given our suspicion that the BDTR model gives more weight to the duration of flow than the other parameters. Long duration of flow creates a more uniform environment in the room and the effect of the direction of airflow is lessened compare to the situations where the duration of flow is limited. Therefore, in the situations with longer duration of flow, the duration of flow becomes the dominant parameter, which is why the BDTR model makes better predictions in those cases. For these reasons we suspect that the BDTR model has not been able to capture the relationships between the airflow direction parameters and the output, but instead has heavily relied on the duration of airflow to the detriment of the other parameters. The unrealistically high value of the coefficient of determination can be accounted for by the outlier effect, since the data points to which the model has shown the best fit are those that have a considerably longer duration of flow compared to the majority of the data points. In contrast to the BDTR model, the GPR model has not ignored the effects of airflow direction by relying heavily on the duration of flow parameter, which is why in Figure 5 (b) the values of absolute error of prediction for the GPR model have not followed the same trend as Figure 5 (a). Further analysis is required to determine if the BDTR model's overall performance is poor, but the represented evidence strongly indicates that the performance of the BDTR model is not as impressive as the R2 values indicate. Given that such a comprehensive evaluation requires more data than available to us, in the interest of time and efficiency we have chosen to use the GPR model instead.

2.3 Results

2.3.1 Dynamics of the Environment

The presented results in this paper are intended to reflect the potentials that the proposed HVAC system has in terms of savings in energy consumption. Figure 6 (a) presents the amount of improvement in the thermal comfort index of the occupant when the HVAC system has been working for a duration of 60 seconds in the CFD Model. By using the trained GPR Model the improvements in thermal comfort index for non-simulated directions have been computed. According to Figure 6 (a) the suitable direction of the airflow for the given duration and environment is one that directs the airflow towards the occupant. Note that the shown energy-efficient airflow direction in Figure 6 (a) is not among the 6 simulated cases (Figure 3), rather it is the prediction of the GPR Model.

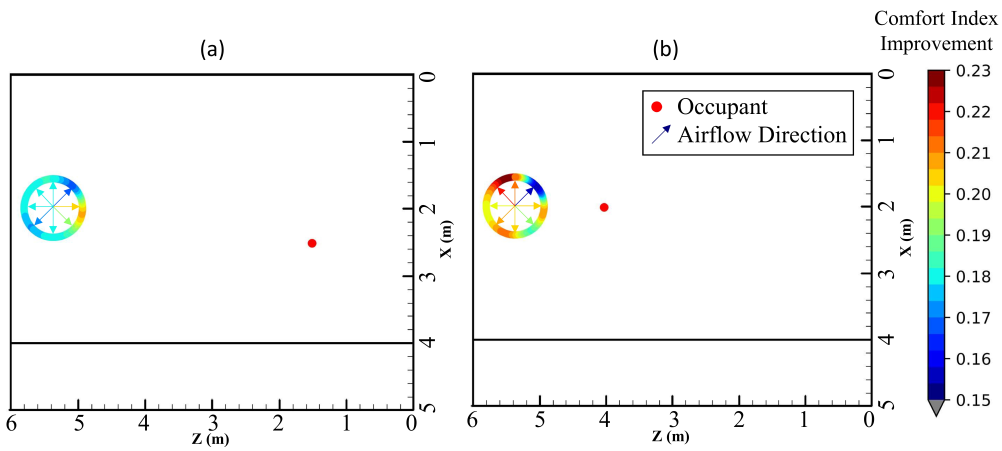

Figure 6: Effect of direction of airflow on changes in the thermal comfort index (Duration=60 seconds)

Figure 6 (b) shows the improvements in the thermal comfort index of an occupant sitting in a different location. Similar to Figure 6 (a) the results belong to the case, where the HVAC system has worked for a duration of 60 seconds. The contrasts between Figure 6 (a) and Figure 6 (b) shed light on challenges that the control system will have to face and overcome. According to Figure 6 (b) in that situation the efficient direction of airflow is not in the direction of the occupant. As mentioned before, in our simulated CFD Model and the trained GPR Model the direction of the airflow has been constrained so that there exists a 30° angle between ceiling and the airflow direction. Thus, if the occupant is sitting close to the diffusor, setting the airflow direction towards the occupant would result in a situation where the airflow passes above the head of the occupant by a distance. This limitation in the flexibility of the simulated HVAC system has made the problem of finding the efficient direction of airflow considerably more complex. The energy-efficient direction of airflow in Figure 6 (b) is one that hits the walls first. Note that despite the geometrical symmetry of the room (with respect to the X=2 line), the improvements in the thermal comfort index are not symmetric. This can be explained by taking into account that one of the walls is in contact with outside warm air which creates a temperature gradient within the room that causes the aforementioned non-symmetric behavior.

As demonstrated by the drawn contrasts between Figure 6 (a) and Figure 6 (b), the optimal behavior of the HVAC system is affected by many parameters such as the airflow dynamics of the room, the imposed temperature constraint and properties related to the occupant. In real-world applications other issues such as barriers (desks, monitors, etc.) can potentially cause further complexity. The control algorithm must be able to learn the optimal behavior despite all of these complexities, while requiring only a minimum number of sensors.

2.3.2 Energy Consumption

In this section, we have examined the potentials for reduction in energy consumption that the proposed flexibility in operation of HVAC system will be able to facilitate. A common practice in HVAC control studies is to define the thermal comfort as a constraint and then minimize the energy consumption of the system. For the purposes of this section, the minimum thermal comfort index will be set to 0.95. The underlying hypothesis is that the energy consumption of a smart, comfort driven HVAC system is sensitive to the variations in the airflow direction at the point of diffusion. Using the CFD simulation and the trained GPR Model, we have computed the duration of airflow necessary for each airflow direction to bring the occupant to the desired thermal comfort index. Since the airflow rate is equal for all directions, for the purposes of this study we will assume that the energy consumption is proportional to the duration of airflow. However, we must be mindful that the relationship between airflow and energy consumption is not necessarily linear in all settings. Figure 7 demonstrates the relationship between the horizontal angle of airflow and the duration of airflow (proportional to energy consumption). The horizontal axis shows the angle between the airflow direction and the direction of the occupant from the diffusor, while the vertical axis shows the duration of airflow required for the occupant to reach the desired thermal comfort index. Positive and negative values on the horizontal axis respectively indicate counter-clock-wise and clock-wise angles between the direction of the occupant and the direction of airflow. The solid line (red) shows the average results for the cases, in which the occupant is sitting at locations where Z=0.5 or 1 or 1.5 (21 locations), with the shaded area representing a variation of one standard deviation from the average. The dash line (blue) belongs to the case where the occupant is sitting at (X=2.5, Z=1.5). In Figure 7 the efficient direction of airflow is the direction that blows the air in the direction of the occupant since this direction has a lower duration of airflow and therefore a lower level of energy consumption. Also note that the consistency of the results for different occupant location as demonstrated by a relatively narrow shaded area corresponding to a smaller standard deviation Figure 8 is similar to Figure 7, except that the occupant is sitting closer to the diffusor at locations where Z=4.5 or 5 or 5.5 for the solid line (red) and the corresponding location to the dash line (blue) is (X=2, Z=4.5). Unlike the previous case, here the efficient angle of airflow is not at the same direction as the occupant's location. Note the more complex behavior and lack of consistency in Figure 8 compared to Figure 7 that is a result of the airflow dynamics of the room and the limitations of the HVAC system.

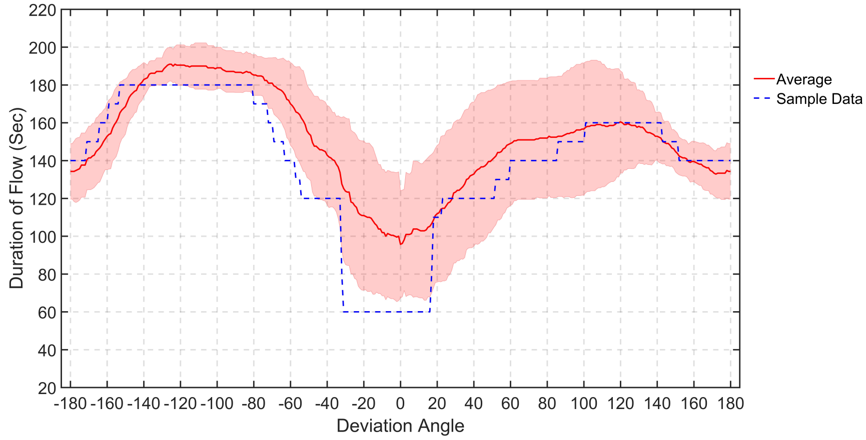

Figure 7: Required duration of airflow in the CFD model that brings the occupant to the desired thermal comfort index versus the deviation angle from the location of the occupant.

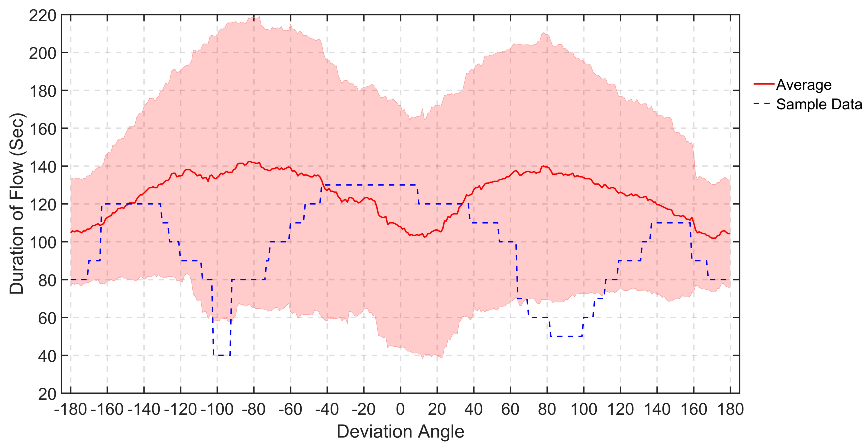

Figure 8: Required duration of airflow in the CFD model that brings the occupant to the desired thermal comfort index versus the deviation angle from the location of the occupant.

To examine the second hypothesis of this paper, namely that significant energy savings can be achieved by using flexible HVAC systems, we will assume a uniform probability distribution for the location of the occupant (occupant has an equal probability of sitting in any location in the room). To account for this assumption, we use the uniformly distributed occupant location presented in Figure 1 and then we average over all user locations. The percentage of energy savings for each direction is calculated in comparison to the uniform diffusion.

$$Energy\ Savings\ for\ direction\ i = 100 \times \left(1 - \frac{E_i}{E_{uniform}}\right) \tag{11}$$

E_i is defined as the amount of Energy needed to bring the occupant to the 0.95 thermal comfort index while diffusing the air in direction i. $E_{uniform}$ is defined as the amount of Energy needed to bring the occupant to the 0.95 thermal comfort index with a uniform diffusion. The average of the results for all the airflow directions (from 1 to 360 degrees by a 1° step size) has been used as an estimate for uniform diffusion. As mentioned before, the airflow rate is constant in our simulations, so the energy consumption is proportional to the duration of airflow. Assuming that a control algorithm can effectively identify the optimal airflow direction, we compute the Maximum Energy Savings for each occupant location.

$$Maximum\ Energy\ Savings = max_i\Big(100 \times \Big(1 - \frac{E_i}{E_{uniform}}\Big)\Big) \tag{12}$$

For each one of the 77 occupant locations the Maximum Energy Savings has been calculated. According to the results, the average amount of energy saved by using the envisioned HVAC system with a competent control algorithm is 59% per cycle. The amount of the saved energy varies between 20% and 92% per cycle, depending on the location of the occupant. The histogram in Figure 9 provides some insight into the potential savings that the proposed system can yield.

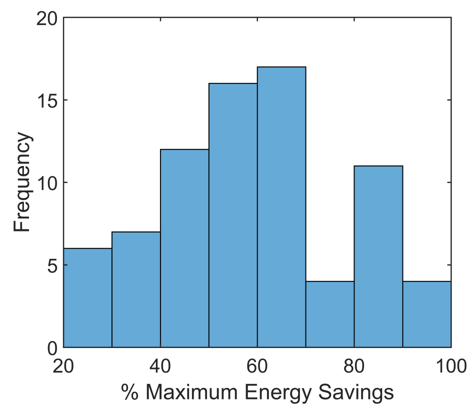

Figure 9: Histogram of the Maximum Energy Savings per cycle achieved in the 77 considered occupant locations.

Figure 10 allows for a better understanding of the relationship between the occupant's location and the potential for savings in energy consumption by using the proposed platform. According to Figure 10, in locations that are closer to the diffusor higher savings can be achieved. This should be taken into account when choosing the location of a diffusor or the number of diffusors in a real-world building. Note that the actual amount of energy savings brought on by the proposed HVAC system also depends on the number of heating/cooling cycles. It is likely that the envisioned HVAC system would have to go through a higher number of operation cycles compared to the case of uniform diffusion.

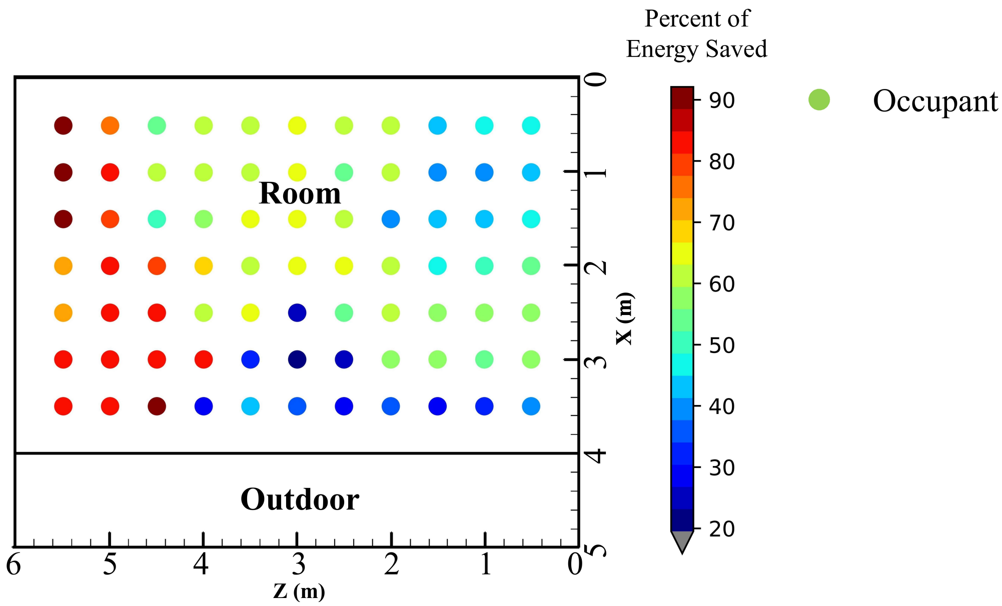

Figure 10: Potential energy savings (per cycle) for each of the considered occupant locations.

2.4 Conclusion

In this paper, we introduced the idea of HVAC systems with more flexibility in terms of possible actions they can take within a thermal zone. A specific case of such HVAC systems, namely an HVAC system that is capable of adjusting the direction of airflow was investigated. A CFD model of the described setting was created and the results of the simulation were further generalized by using Gaussian Process Regression method to alleviate the need to perform a large number of computationally expensive CFD Simulations. The results indicate that the idea of enabling the HVAC system to have control over the direction of airflow at the point of diffusion, offers a very promising prospect. The average energy savings of 59% per cycle (assumed to be proportional to the airflow duration) in comparison to a uniform-diffusion HVAC system calls for further research into this potentially effective and underexplored area of research. Notwithstanding the demonstrated promise of the proposed ideas, the results show that a successful real-world design and implementation of such systems is predicated on effectively addressing the complexities and barrier that lie in the way. One of the challenges that such a system will face is the issue of designing an effective control algorithm. It was discussed in the previous sections that the optimal response at each state, as well as the potential for energy savings dramatically varies with parameters related to the occupant's location, the limitations of the HVAC system and the airflow dynamics of the thermal zone. Taking into account the infeasibility of installing a high number of sensors, the control algorithm must be capable of learning how to effectively control the HVAC system while using only a minimal amount of information. The envisioned HVAC system should also address the ventilation requirements. Our future line of research will focus on creating an interactive framework, in which an intelligent control algorithm would learn how to control the HVAC system. We are also in the process of designing and executing relevant empirical studies.

2.5 Acknowledgment

This material is based upon work supported by the Institute for Critical Technology and Applied Science (ICTAS) at Virginia Tech. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ICTAS.

3 Research Study 2: Human-in-the-loop Energy-Efficient HAVC Operations through Active Directional Flow using Reinforcement Learning

Authors: Milad Abedi[1], Farrokh Jazizadeh[1], Bert Huang[2], Francine Battaglia[3]

[1]Via Department of Civil and Environmental Eng., Virginia Tech, Blacksburg, VA, 24061, USA
[2]Department of Computer Science, Virginia Tech, Blacksburg, VA, 24061, USA
[3]Department of Mechanical & Aerospace Engineering, University at Buffalo, NY, 14260,USA

Key Words: HVAC Control, Directional Airflow, Thermal Comfort, Reinforcement Learning, User-Centric HVAC systems

3.1 Introduction

Enhanced control of Heating, Ventilation and Air-Conditioning (HVAC) systems has been the subject of numerous research efforts. Improvements in the operation of HVAC systems have been sought in two areas of HVAC energy consumption and occupant thermal comfort. In the United States, buildings consume 39% [1] of the total energy, with the HVAC system being responsible for 48% [3] and 44% [2] of the energy consumption in residential and commercial buildings respectively, resulting in high incentives for energy efficiency improvements.

In the conventional control methods, the HVAC systems are configured to conform to the temperature requirements of a space (measured at the location of a thermostat), assuming that the temperature set-point of the thermostat is a sufficiently accurate indicator of the thermal comfort tendencies of the occupants. However, this assumption is not entirely accurate for a number of reasons. Firstly, the thermal comfort state of the human subject often depends on many parameters other than the surrounding air temperature, such as airflow velocity, clothing, and metabolic rate [30]. Secondly, there is no guarantee that the temperature at the arbitrary location of the thermostat is the same as the air temperature in the vicinity of the occupant. One of the approaches taken to tackle the deficiencies in the conventional approach to HVAC control is to account for personalized and contextual thermal comfort state of the occupants in driving thermostat set points [4–6]. While this strategy helps improve upon the conventional control approach, there could be room for improvement. For instance, in a house with one thermal zone consisting of multiple rooms and a single thermostat, the response of the HVAC system to changing the temperature set-point of the thermostat would be to direct the conditioned air to all the zones, occupied and unoccupied alike until the temperature at the location of the thermostat reaches a certain range. Logically, conditioning the air in unoccupied zones serves no purpose and is an unnecessary energy expenditure.

An implemented solution to this shortcoming is to divide larger areas with a single air conditioning system (i.e., a larger thermal zone) to a number of smaller zones with individual thermostats and diffusors capable of binary operations (i.e., opening and closing). By augmenting these systems with occupancy detection sensors, the HVAC system could avoid excessive use of energy by conditioning the air only in occupied areas. However, these systems are also subject to a number of limitations. As is the case in the conventional approach to HVAC control, the zoning solution also suffers inaccuracies from assuming that the thermal condition at the location of the thermostat is representative of the condition at the location of the occupant. In reality the temperature might not be uniform across the room and response of the HVAC system might not create a uniform change. Abanto et al. [31] have demonstrated that the operation of the HVAC system in a particular room equipped with uniform diffusors does not necessarily lead to a uniform temperature distribution within the room.

In order to tackle the non-uniformity of temperature distribution we have investigated the advantages of replacing the traditional uniform diffusors with devices that would enable the HVAC system to control the direction of airflow at the point of the air supply. This approach would allow the HVAC system to build a local climate surrounding the occupant (Figure 11) thereby obviating the need to condition the air in the entire room. This concept could potentially improve energy efficiency and personalized thermal comfort in the environment.

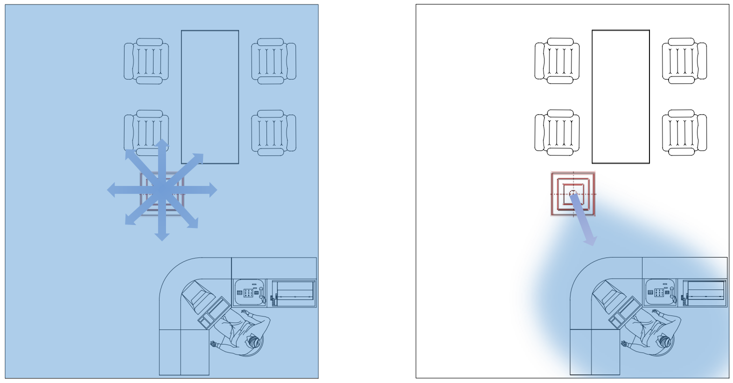

Figure 11: Uniform diffusion vs targeted airflow

We have demonstrated that enabling the HVAC system to control the direction of the airflow at the location of the air-vent could lead to energy savings [32]. However, the success of this solution is predicated on the ability of the HVAC system to learning the optimal action strategy, that is the HVAC system must be able to determine the best direction of airflow at the location of the air-vent. In this study, we have investigated the characteristics and feasibility of a control framework that enables the HVAC systems to autonomously determine the best airflow direction by accounting for energy consumption and occupant thermal comfort.

A control framework based on the reinforcement learning (RL) approach, which is an artificial intelligence (AI) learning method enabling a system to operate as an autonomous agent, will be evaluated in this work. The proposed framework will aim to simultaneously maximize the occupant thermal comfort while minimizing the energy consumption. It was shown in a previous study [32] that the location of the occupant within the room was an important piece of information to determine the best direction of airflow. However, for the proposed framework herein, we sought to avoid introducing any knowledge regarding the location of the occupant into the control framework. This decision was made in order to address the privacy concerns that would arise from using effective indoor localization technologies such as vision-based techniques.

The paper has been organized as follows. A review of the research efforts on user-centric control in HVAC systems has been presented in Section 3.2. Section 3.3 presents the control framework including the rationale for the choice of control algorithm, and different components of the controller. In Section 3.4, the performance of the control framework has been evaluated. Finally, Section 3.5 concludes the paper and presents future directions of our research.

3.2 Background Review

As a part of research efforts to enhance the operations of the HVAC systems, some researchers have focused on designing and building HVAC control frameworks in which the occupant is an integral part of the control process. The inclusion of the occupant in the control framework is achieved through a multitude of approaches. The following is a concise discussion of the current state of the literature in human-in-the-loop HVAC control. In order to achieve a user-centric HVAC control, the framework must be capable of understating the needs of the user and be capable of addressing those needs.

A considerable portion of the literature on user-centric HVAC control uses the predicted mean vote (PMV) model [30] for evaluation of the thermal comfort state of the occupants [6, 33]. However the PMV model has been designed to reflect the average thermal comfort tendencies of a group of people rather than that of an individual. Therefore, a number of studies have endeavored to build personalized frameworks for occupant thermal comfort evaluation so as to render the HVAC control more occupant-centric. Among these approaches are the participatory sensing frameworks (e.g. [4, 10, 13, 15, 17, 34–36]) in which, the occupant is asked to provide the system with feedback regarding their thermal comfort state and the feedback is then used to construct the control signal. In order to build a less intrusive and more personalized control framework, other researchers [12, 37–41] have used physiological measurements to evaluate the thermal comfort state of the occupants. All of these efforts have been made to tackle the first requirement of a user-centric HVAC control framework, i.e. the ability to understand the needs of the user (demand).

On the control side, one line of research is aimed at integration of user-centric operations with minimal intrusion to the internal operations of the HVAC systems. Thus, these research efforts have aimed to offer solutions for determination of the thermostat temperature set-point that would take the thermal comfort state of the occupant into account. For example, Bermejo et al. [6] have used fuzzy logic controllers to modify the temperature set-point of the thermostat to optimize the thermal comfort state of the occupant as defined by the PMV model [30]. In another instance, Nabil and Samir [5] have proposed a temperature set-point controller that attempts to optimize the energy consumption of the HVAC system, while accounting for the thermal comfort of the occupant via a predefined allowable temperature range.

Other studies [33, 42–44] have tried to further modify the internal operations of the HVAC system thereby achieving a level of control beyond adjusting the temperature set-point of the thermostat. For instance, in an optimal control framework Nassif et al. [33] have extended the level of control to include such components as "supply air temperature set-point", and "supply duct static pressure set-point" while the cost function is defined as a two-objective function comprising of the occupant comfort and energy consumption components. The occupant thermal comfort has been defined as the "predicted percentage of dissatisfied", an index that is defined base on the PMV model.

A number of studies have explored the potential benefits that might be achieved via augmenting the existing HVAC systems with extra devices and sensors with minimal or no modification in the inner components of the HVAC system. Lu et al. [45], Dong and Andrews [46], Dong and Lam [47] used a network of sensors to detect the occupancy state of the zones and to recognize occupancy patterns so as to use this information to turn the HVAC system on and off. The periodic nature of occupancy patterns allows for building of models that provide predictive knowledge of future conditions. This has enabled researchers to successfully use model predictive control approaches, which allows them to incorporate the foresight arising from the predictability of occupancy patterns in HVAC control[45–51].

Contrasting to this approach, a number of studies have proposed modifications to the HVAC system that would enable the existing systems to control the airflow to each room thereby resulting in a more fine-grained level of control [52]. In conjunction with the network of sensors that add the ability to detect and predict the occupancy state of each room, it is possible for the HVAC system to focus its operation only in occupied areas. For instance, Feldmeier and Paradiso [9] developed a network of sensors and actuation hardware that allowed the HVAC system to locate the occupants in the building, detect their thermal comfort state via direct measurement of related parameters on the skin, and direct the air only to the occupied areas.

In this study, we have sought to create further adaptation capacities in the HVAC system by enabling control over the direction of airflow at the point of the ceiling air-vent. In so doing, a considerable increase would be achieved in the level of fine control capabilities of the HVAC system compared to other solutions. Permitting the HVAC system the ability to control the direction of airflow, the control framework will no longer be subject to the limitations of uniform diffusion paradigm. By virtue of the new action capabilities (i.e., control of airflow direction), the system can forgo the assumption of a uniform temperature distribution within the room and incorporate the variations in temperature distribution within the room and its effect on the occupant. With the non-uniformity of the room environment, the control over the direction of airflow could be utilized to leverage the non-uniform nature of the environment towards a more pleasant and energy efficient user experience. For this objective to be realized, a control framework capable of addressing the complexities of the environment must be developed. Moreover, the framework must be able to address the privacy concerns surrounding the indoor localization technology, by being capable of operating in the absence of knowledge regarding the location of the occupant.

In view of the above-mentioned requirements, we have chosen RL to serve as the control algorithm in charge of driving the direction of airflow. In context of RL, the control algorithm will learn based on the information gathered through interactions with the environment where the agent will receive feedback (reward) from the environment with regard to the favorability of the action taken. Through iterations of this interaction the agent will learn the appropriate behaviors, which is the action policy that leads to the maximum expected overall reward.

In recent years, researchers in numerous fields have turned to RL to address their needs for reliable, autonomous control algorithm that could handle the complexities of the control process. The intuitive nature of the notion of reward function, and the obviation of the need to have prior models of environmental processes has also encouraged researchers in the field of building systems control to utilize RL. For instance, in a simulation-based study, Chen et al. [53] used RL to enable autonomous control of the HVAC and window systems where the RL agent is able to turn the HVAC system on and off in addition to being able to open and close the windows. The reward function has been designed to penalize against energy consumption and deviation from the target temperature and humidity. In another simulation based study, Baghaee and Ulusoy [54] have used RL to control the "ventilation rate" of the HVAC system. The reward function in this study was designed so as to penalize against deviations from the desired CO_2 content and energy consumption. Wei et al. [55] have investigated the feasibility of utilizing Reinforcement Learning for control of the HVAC system in a simulated environment. In their study, the agent is capable of choosing the airflow rate to each zone. Their designed reward function penalizes against the cost and deviations from the temperature set-point. In another simulation based study, Fazenda et al. [56] have used RL to control the temperature set-point of the HVAC system in accordance with the occupancy-related behaviour of the occupants so as to improve comfort of the occupants as well as the energy consumption of the HVAC system.

In this study, we have evaluated the potential of Reinforcement Learning to serve as the autonomous control algorithm that will guide the expanded adaptive capacity created by enabling the HVAC system to have control over the direction of airflow. The performance of the RL algorithm has been evaluated in terms of its ability in operating subject to a dearth of information regarding the location of the occupant, brought on by privacy concerns. In order to evaluate the performance of the RL algorithm under real-world training-time limitations, we have introduced constraints on RL training time. Moreover, the intuitive design of the reward function makes it easy for the end-users to convey their comfort/cost expectations to the control framework to enable a more personalized control objective.

3.3 Reinforcement Learning as the Control Algorithm

Selection of the Reinforcement Learning algorithm as the control algorithm in charge of adjusting the direction of airflow, was based on certain contextual requirements. Firstly, given the fact that the location of the occupant is a major factor in determining the best direction of airflow at the air-vent, knowledge of the relative location of the occupant can be an invaluable piece of information for the control algorithm. However, utilization of indoor localization technology will introduce a number of privacy concerns. In order to address privacy concerns, the control algorithm must be capable of performing in the absence of information about the location of the occupant. Secondly, a number of thermal zone properties such as room dimensions and location of the air-vent can have a major influence in determining the best direction of airflow. However, in order for the control algorithm to be capable of performing in a variety of environmental settings, the control algorithm must be capable of performing in the absence of prior information about the properties of the environment. In other words, the control algorithm must be able to learn through exploration of the unknown environment. These considerations have driven us towards choosing RL to serve as the control algorithm. The following is a discussion of components of the RL control framework.

3.3.1 State

In order for the control algorithm to choose an action, first it must have an understanding of the current situation. The knowledge of the current situation is given to the RL agent via the definition of the State. State, is the information given to the agent (often an array of numbers) regarding the environment in so far as it relates to the operations of the RL agent. In this study, we have defined the state to be the temperature at the location of the occupant. In real-world scenarios the ambient temperature at the location of the occupant can be easily measured given the ubiquity of various types of temperature sensors.

3.3.2 Action

In each interaction with the environment, the agent is capable of choosing the direction of airflow, represented by a unit vector of length 3 corresponding to X,Y, and Z axes. In addition to the direction of flow, the agent is also capable of choosing the duration of flow for the interaction in question. Thus in each interaction the agent will determine the element values for an array of length 4 where the first 3 elements represent the unit direction vector and the 4^{th} element determines the duration of flow.

3.3.3 Reward

In the context of Reinforcement Learning, Reward is a scalar that represents the feedback from the environment to the actions taken by the agent. The agent will receive higher rewards for favorable actions and by contrast it will receive lower rewards for unfavorable actions. Definition of the rewards is one of the most challenging steps in designing an RL controller as any error in the definition of the reward function will inevitably lead to error in the performance of the RL agent. Also the definition of the reward could have a significant effect on the training time of the RL agent.

To define the reward in the context of HVAC system control, we have examined our expectations from the control agent and constructed the definition of the reward so as to meet those expectations. As discussed in the Introduction section, the ideal control algorithm should be able to improve the thermal comfort state of the occupant while using a minimum amount of energy. In view of this objective, we have defined the reward as follows:

$$Reward = a_1 \times \Delta_{TC} - a_2 \times E \tag{13}$$

Where:

Δ_{TC} = Change in the thermal comfort state of the occupant
E = Energy consumed for taking the action
a_1, a_2 = Positive constants

By virtue of the definition of the rewards as shown in Equation 13 the agent will receive higher rewards for actions that cause greater improvements in the thermal comfort state of the occupant(s) and it will be penalized for the amount of energy it has consumed to take that particular action. The constant coefficient of a_1 and a_2 represent the relative importance of the energy consumption versus the occupant thermal comfort. Choosing large values for a_1 and small values for a_2 would indicate that the occupant is more concerned about the thermal comfort and less concerned with the energy consumption. By contrast small values for a_1 and large values for a_2 would indicate that the occupant is more concerned about the energy consumption than occupant thermal comfort. Thus the values of a_1 and a_2 are a matter of personal preference and can be effected by considerations such as energy pricing. This explicit definition will allow the end-users to intuitively convey their preferences to the control framework.

3.3.4 Q-Function

In a given state (e.g. $s_i : T_0 = 26\,^{\circ}\mathrm{C}$) the agent should choose an action among all possible actions such that the expected overall reward is maximized (see Equation 14). Note that future expected rewards are discounted according to their latency.

$$a_i = argmax\ _a\Big(E\Big[\sum_{t=T}^{\infty}\gamma^{t-T}r_{_T} \mid S = s_i, A = a\Big]\Big) \tag{14}$$

Where:

A = Action
S = State
γ = Discount factor for future rewards
r_T = Reward at time step T

In evaluating Equation 14, the agent will compare look for and choose the action that will result in the largest value for the expected overall reward. The expected overall reward given a particular state and action is called the *Action-Value* and is denoted by $Q_{(s,a)}$.

$$Q_{(s,a)} = E\Big[\sum_{t=T}^{\infty} \gamma^{t-T} r_T \mid S = s_i, A = a\Big] \tag{15}$$

In cases where the space of possible State-Action pairs is discrete and small, one could use a table to store the values of $Q_{(s,a)}$ for each pair of s, a. However, as the number of possible State-Action pairs increases, the tabular representation for $Q_{(s,a)}$ becomes more and more infeasible. In the case of controlling the direction of airflow, both the State space and the Action space are continuous and therefore an infinite number of State-Action pairs can be defined. In cases such as this a common approach (Function Approximation) is to use a continuous function to represent $Q_{(s,a)}$, and update the model parameters as more information is observed (Equation 16).

$$Q_{(s,a)} \approx \hat{Q}_{(s,a)} = f_{(s,a|\theta)} \tag{16}$$

Where:

$\hat{Q}_{(s,a)}$ = Approximation of $Q_{(s,a)}$
θ = Model parameters

In choosing the model to represent $\hat{Q}_{(s,a)}$ one must be careful not to impose unwarranted restrictions on the Action-Value space. For instance if the $Q_{(s,a)}$ has a highly non-linear shape, choosing a linear model as the approximating model can lead to erroneous outcomes. Perhaps the most challenging task in choosing an appropriate model for the Action-Value function is the lack of prior information about the particular $Q_{(s,a)}$ in question at time. In the absence of prior knowledge about the characteristics of $Q_{(s,a)}$ choice of models that can mimic a wide variety of shapes can help avoid the aforementioned issues. In this study, we have chosen a Gaussian Process Regression model to approximate $Q_{(s,a)}$. Gaussian Process Regression is a non-parametric regression model whose shape is determined to a great extent based on the observed data rather than prior parameterization assumptions.

3.3.5 The Learning Algorithm

In the context of Q-Learning with Function Approximation, the task of learning is the training of the $\hat{Q}_{(s,a)}$. The following is the derivation of an iterative algorithm for updating the $\hat{Q}_{(s,a)}$ function. Per it's definition, Q function is the expected total discounted reward if the agent takes action a_T in the current state s_T (Equation 17).

$$Q_{(s_T,a_T)} = E\Big[\, r_T + \sum_{t=T+1}^{\infty} \gamma^{t-T} r_t \Big] = E\Big[\, r_T \Big] + \gamma \cdot E\Big[\sum_{t=T+1}^{\infty} \gamma^{t-(T+1)} r_t \Big] \tag{17}$$

After taking the action a_T in the current state the reward r_T will be received, and substituting $T' = T + 1$ in Equation 17 will result in:

$$Q_{(s_T, a_T)} = \; r_T \; + \; \gamma \cdot E\Big[\sum_{t=T'}^{\infty} \gamma^{t-T'} r_t\Big] \tag{18}$$

Using the definition of the Q function, Equation 18 can be rewritten as:

$$Q_{(s_T, a_T)} = \; r_T \; + \; \gamma \cdot Q_{(s_{T'}, a_{T'})} \tag{19}$$

Given that the behaviour policy of the system is to take the action which maximizes the value of the Q function, Equation 19 can be rewritten as:

$$Q_{(s_T, a_T)} = r_T + \gamma \cdot \max_{a_{T'}} Q_{(s_{T'}, a_{T'})} \tag{20}$$

Equation 20 is then used to develop the following iterative update algorithm to train the $\hat{Q}_{(s,a)}$ model.

Algorithm 1 Training the Q function

1: **while** $\Delta \geq \underline{Tolerance}$ **do**
2: $\quad Y \leftarrow r_T + \gamma \cdot \max_{a_{T'}} Q_{(s_{T'}, a_{T'})}$
3: $\quad \Delta \leftarrow \Big\| Y - \hat{Q}_{(s_T, a_T)} \Big\|$
4: $\quad$ Train $\hat{Q}$ via regression ($Inputs = State - Action\ pairs;\ Output = Y$)
5: **end while**

3.3.6 Exploration vs Exploitation

In the context of Q-Learning, at each state the RL agent could choose the action that maximizes the value of $\hat{Q}_{(s_T, a_T)}$ (acting greedily, also called Exploitation). However if the agent were to always act greedily, there is a danger that the agent might not sufficiently explore the possible range of actions. In order to force the agent to sufficiently explore the environment as apposed to repeating a particular routine it deems to be best based on prior experiences, we will force the agent to act randomly every now and then. In order to strike a reasonable balance between exploration and exploitation, we will use an algorithm called the $\epsilon - Greedy$ algorithm. The $\epsilon - Greedy$ algorithm will force the agent to act randomly at each step with the probability of ϵ.

$$Action = \begin{cases} Greedily & Probability = 1 - \epsilon \\ Randomly & Probability = \epsilon \end{cases}$$

As the time passes and agent becomes more experienced with a better understanding of the environment, the unexplored portions of the State-Action space become smaller, and therefore there is less need for the agent to act randomly. For this reason, at each time step we will decrease the probability of acting randomly. Accordingly, at time step t, the probability of acting randomly has been defined in Equation 21:

$$\epsilon = \frac{1}{t} \tag{21}$$

3.4 Results

In order to evaluate the performance of the developed control framework, we have designed a synthetic environment in which to test the capabilities of the RL algorithm in controlling the direction of airflow. To do so, the results of a CFD model of a typical room (Figure 12) with a single diffusor and a single exhaust, developed in the commercial software ANSYS Fluent, were used. The details of the simulation set-up and model properties can be found in [32].

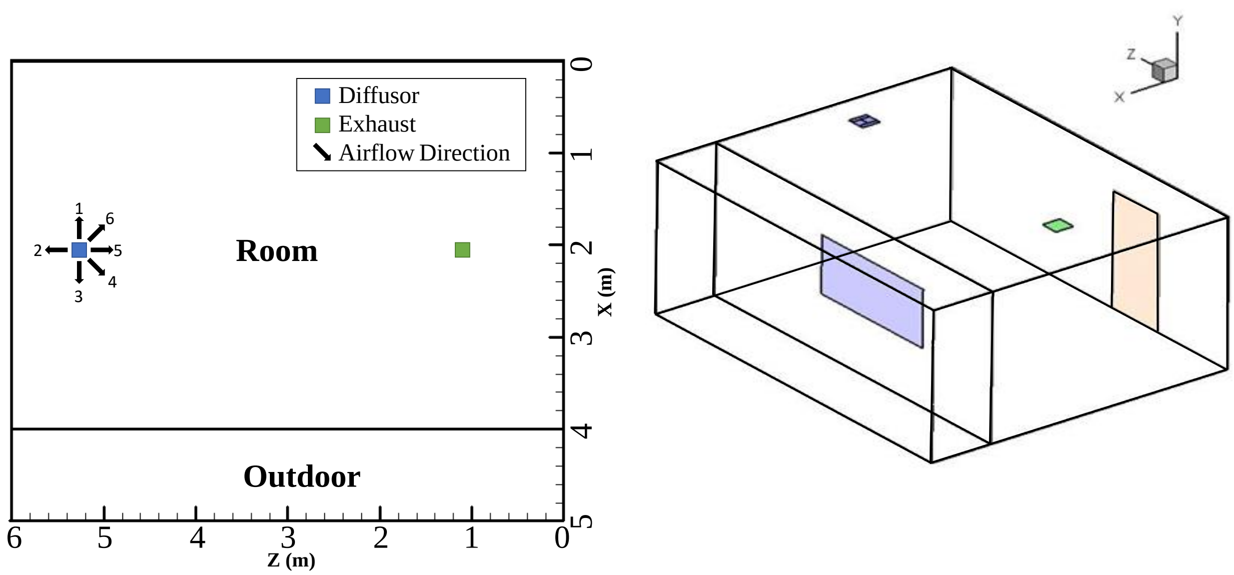

(a) Plan View (b) 3-D View

Figure 12: The room simulated in the CFD model

Due to high computational expense of running the CFD model, simulating airflow for each direction of interest via the CFD model is not feasible. For this reason, results of CFD simulations for 6 airflow directions (Figure 12(a)) were used as a bases to estimate the results for other directions of airflow. We have used a Gaussian Process Regression (GPR) model to estimate the results of the CFD analysis for non-simulated airflow directions [32]. To build the simplified model of the thermal environment, the amount of change in the temperature at the location of the occupant has been expressed in the following form:

$$\Delta T_{oc} = F_{Dur} \times F_{Dir} \tag{22}$$

Where:

ΔT_{oc} = Change in the temperature at the location of the occupant
F_{Dur} = Duration of Flow Factor
F_{Dir} = Direction of Flow Factor

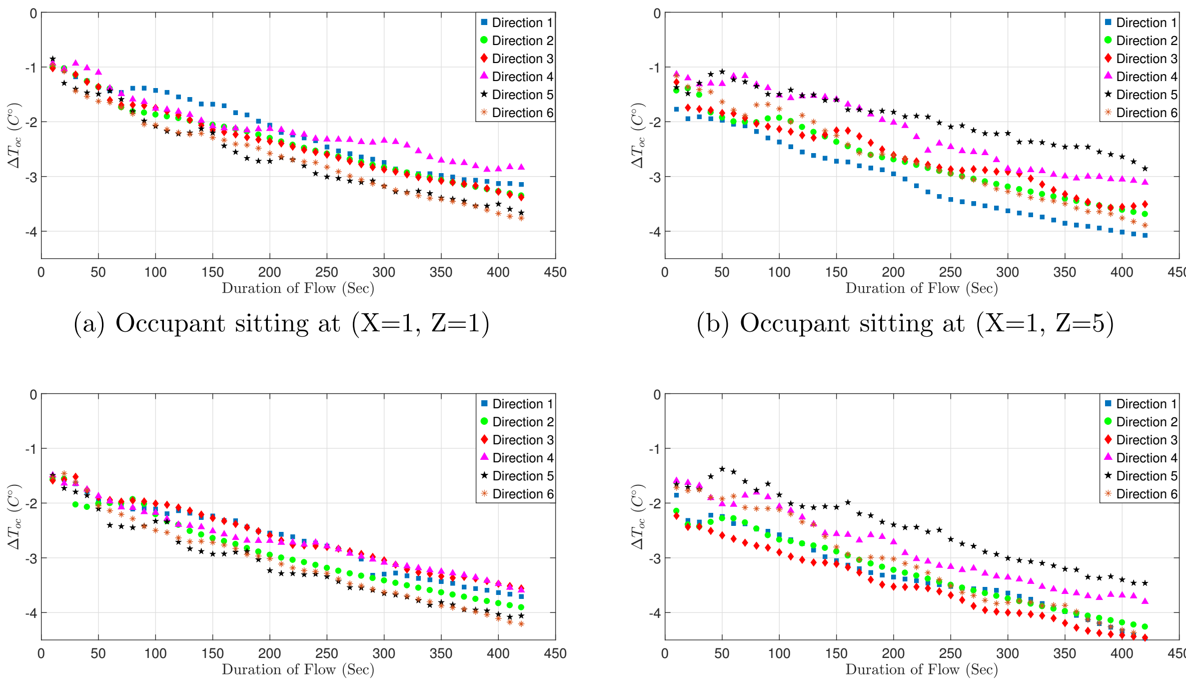

(a) Occupant sitting at (X=1, Z=1)

(b) Occupant sitting at (X=1, Z=5)

(c) Occupant sitting at (X=3, Z=1)

(d) Occupant sitting at (X=3, Z=5)

Figure 13: The relationship between the Duration of Flow and the Change in Temperature at the location of the occupant with different airflow directions

The choice of this model was inspired by the observation that the relationship between the change of temperature at the location of the occupant and the duration of airflow has the same linear behaviour for different directions of airflow. In Figure 13 the amount of change in temperature at the location of the occupant (for 4 different occupant locations) versus the duration of flow for the 6 simulated airflow directions has been plotted. As can be seen in Figure 13, the quasi-linear nature of the relationship between the duration of flow and the ΔT_{oc} does not change with the change in the direction of airflow.

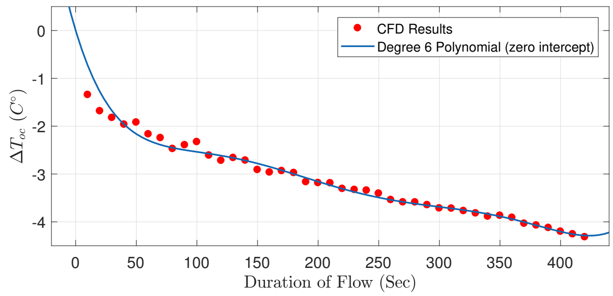

Figure 14: The relationship between Duration of Flow and the Change in Temperature at the location of the occupant using the most effective direction of airflow (the direction causing the fastest change in temperature)

For a given location of the occupant, in order to build the simplified environment, the Duration Factor (F_{Dur}) model will be built by measuring the temperature drop for different durations of flow. To do so, of all the airflow directions, the direction that causes the greatest amount of change in the temperature at the location of the occupant will be selected. For this airflow direction, the Duration of Flow vs ΔT_{oc} relationship will be modeled via a degree 6 zero-intercept polynomial regression model as shown in Figure 14. The resulting regression model will be used as the Duration Factor (F_{Dur}). Thus, the duration factor represents the amount of change in the temperature at the location of the occupant for different durations of airflow when the direction of airflow has been chosen such that it creates the maximum change in the temperature . For other directions of airflow, the temperature change will be calculated by multiplying the Duration Factor by a coefficient that captures the effect of airflow direction (i.e. the Direction Factor).

The model representing the Direction of Flow Factor F_{Dir} will be built based on the results of the GPR model that we developed to estimate non-simulated airflow directions. Using this model we calculated a coefficient that will reflect the effect of direction of flow as shown in Equation 22. To do so we will use the GPR model to measure the amount of temperature change at the location of the occupant for different directions of airflow and we will call it $g_{(\theta)}$ ($theta$ represents the direction of flow). To obtain F_{Dir} we will normalize $g_{(\theta)}$ such that its maximum value equals 1 (Equation 23), corresponding to the most energy efficient direction of airflow. An example of the F_{Dir} function has been presented in in Figure 15 for an occupant sitting in (X=1, Z=1).

$$F_{Dir} = \frac{g_{(\theta)}}{\max_{\theta}(g_{(\theta)})} \tag{23}$$

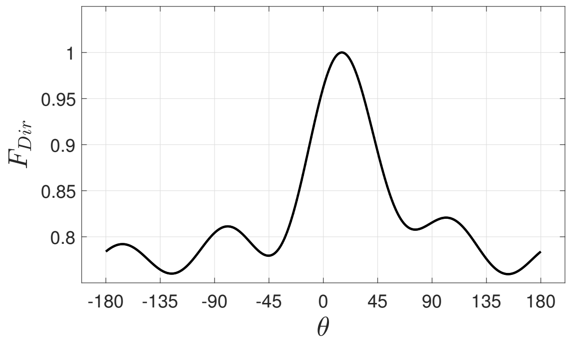

Figure 15: Direction Factor for occupant sitting at (X=1,Z=1)

As for the representation of the human in the environment we have chosen to use the model of the thermal comfort offered by Daum et al. [18] which defines the relationship between the thermal comfort state of the occupant and the surrounding air temperature (Figure 16(a)).

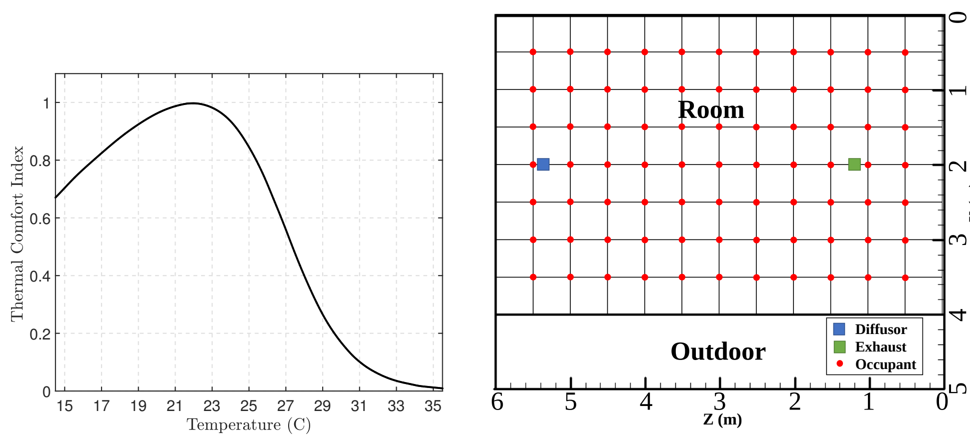

(a) Thermal comfort index [18]   (b) Considered occupant locations

Figure 16: Description of the occupant as a part of the environment

In order to evaluate the potential savings that could be achieved by enabling the HVAC system to control the direction of the airflow, we have considered a total of 77 occupant locations uniformly distributed within the room as shown in Figure 16(b). Starting from the initial temperature of 26 °C, we have calculated the duration of flow required to bring the occupant to the thermal comfort index of 0.95 (Figure 16(a)) for the optimal direction of airflow (the direction of airflow that requires the shortest duration of flow) and for the traditional uniform diffusion, calling them D_{min} & D_u respectively.

In their 2014 paper Ghahramani et al. [4] demonstrated the existence of a linear relationship between the airflow rate and the energy consumption. Under the assumption of a linear relationship between the rate of airflow and the energy consumption, equation 24 estimates the percentage of reduction in the required duration of flow for the directional airflow in comparison to uniform airflow. Given the constant rate of the airflow in our simulated environment, the percentage of reduction in the required duration of airflow is approximately equal to the the percentage of energy saved for the individual diffusor in question.

$$\Delta E \approx 100 \times \left(\frac{D_u - D_{min}}{D_u}\right) \tag{24}$$

In Figure 17(a) we have shown the potential for energy saving that can be reached by enabling the HVAC system to control the direction of airflow for each of the 77 considered occupant locations. As can be seen in Figure 17(a) the location of the occupant has a major influence on the potential for energy saving. Also, occupant locations that are closer to the diffusor have a higher energy saving potential since a smaller portion of the room needs to be cooled in order to target the occupant. Moreover, in occupant locations that are closer to the warmer wall (the wall in contact with the hot outdoor environment) the potential for energy saving is lower than the locations that are further away from the warm wall. In Figure 17(b) the histogram of the values calculated by equation 24 for the potential energy saving save been presented. On average energy saving of %59 per cycle was achieved in the analysis of 77 occupant locations.

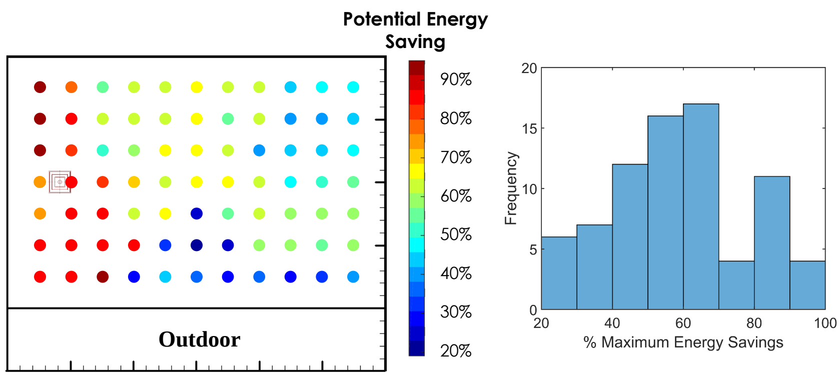

(a) (b)

Figure 17: Energy saving potential analysis

Studying the relationship between the direction of airflow and the duration of flow required to bring the occupant to the desired thermal comfort level (thermal comfort index of 0.95) can shed light onto the dynamics of the environment. To that end we have selected two sides of the room in which to study the relationship between the direction of airflow and the required duration of flow to reach thermal comfort targets (Figure 18). The results of this analysis have been presented in Figure 19(a) and Figure 19(b).

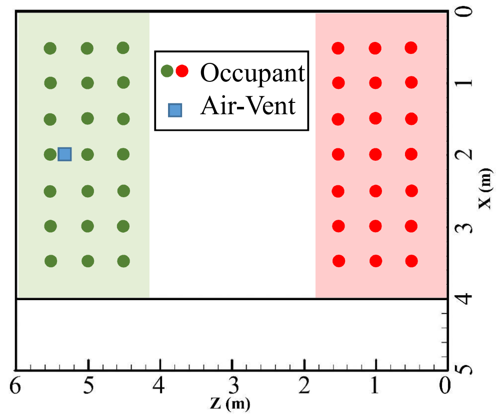

Figure 18: Energy saving potential analysis

In Figure 19(b) the duration of flow required to bring the occupant to the comfort level of 0.95 has been plotted against the direction of airflow (as defined in Figure 20). The red line in Figure 19(b) is the average value of the required duration for 21 occupant locations on the right hand side of the room as show in Figure 18. The shaded area in Figure 19(b) represents one standard deviation from the mean. As reflected in the narrowness of the shaded area, the dynamics in the right hand side are consistent and simple, where the most energy efficient direction of airflow is one that directs the flow towards the occupant.

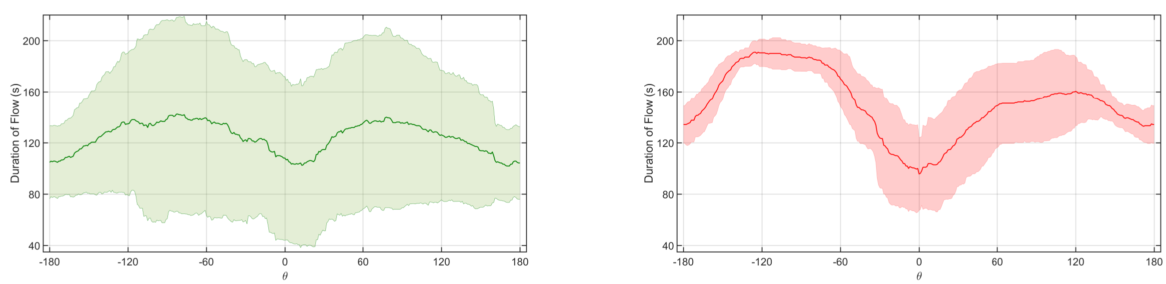

(a) (b)

Figure 19: Energy saving potential analysis

In contrast to the right-hand side, the dynamics of the environment in the left-hand side are more complex as has been reflected in the wider shaded area (Figure 19(a)). The Reinforcement Learning control agent must be capable of dealing with such complexities in the room that could arise from the occupant's choice of location or from the thermal constraints inherent in the environment such as having a wall in contact with the hot outside environment. In order to evaluate the performance of the RL control agent under the described complexities, we have considered four occupant locations for our experiments as shows in Figure 20. The chosen locations vary both in terms of their location characteristics as shown in Figure 19, and their exposure to the thermal constraints of the environment as reflected in their distances from the wall in contact with the hot outside environment.

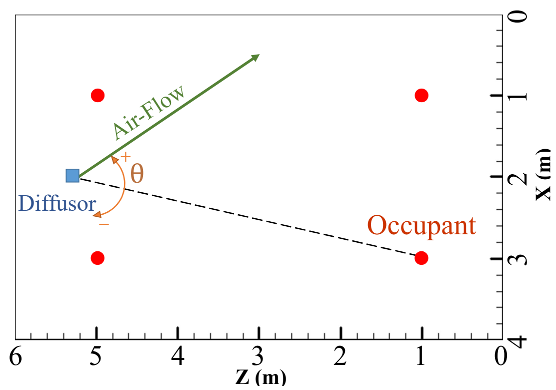

Figure 20: Occupant locations considered in evaluation of the control algorithm

In order to investigate the performance of the Reinforcement Learning agent in controlling the direction of the airflow at the point of the air-vent the RL agent was permitted to interact with the Environment for a total of two cooling cycles. In each cycle, the RL agent was allowed to change the direction of airflow 5 times, while observing the effect of each action after taking it. The rationale behind restricting the length of the agent's learning experience before performance evaluation is to simulate a real world scenario where the directional airflow has been enabled in a room. It is imperative that the controller be able to demonstrate acceptable performance in a short amount of time after installation. Also the agent must be able to quickly address any changes in the thermal environment of a room such as opening of a window. In order for the agent to be able to react briskly to such changes, the learning time for the agent should be as brief as possible.

At the end of the two cooling cycles, the agent has been trained based on a total of 10 experiences (5 interactions per cycle). At this point we made an inquiry about the RL agent's judgment regarding the optimal direction of airflow at the beginning of the day when the air temperature surrounding the occupant (S_0) was approximately 26 °C. Since from the perspective of the RL agent the optimal action is the one that incurs the greatest $Q_{(s,a)}$ value, the response of the agent to the inquiry would be as follows:

$$a_{_{opt}} = arg \max_{a} Q_{(S_0,a)} \tag{25}$$

We have repeated this process for 1000 times and then drew the histogram of the agent's understanding of the best direction of airflow, denoted by θ_{opt} (defined similar to θ in Figure 20). In Figure 21(a) you can see the histogram of the θ_{opt} for the occupant sitting at (X=1, Z=1). Recall that based on the control problem set-up, we expect the optimal policy to maximize the thermal comfort state of the occupant while minimizing the energy consumption. This objective can be reached by selecting the direction of airflow that causes the greatest change in the temperature at the location of the occupant for a given duration of airflow (i.e. the most energy efficient airflow direction). In construction of the Equation 22 we demonstrated that the effectiveness of every airflow direction in changing the temperature at the location of the occupant is accounted for by using a direction of flow factor (F_{Dir}) where the more energy effective direction is the one with a greater value of F_{Dir}. It follows that the optimal action policy is one that aims to choose directions of airflow with the maximum F_{Dir} value.

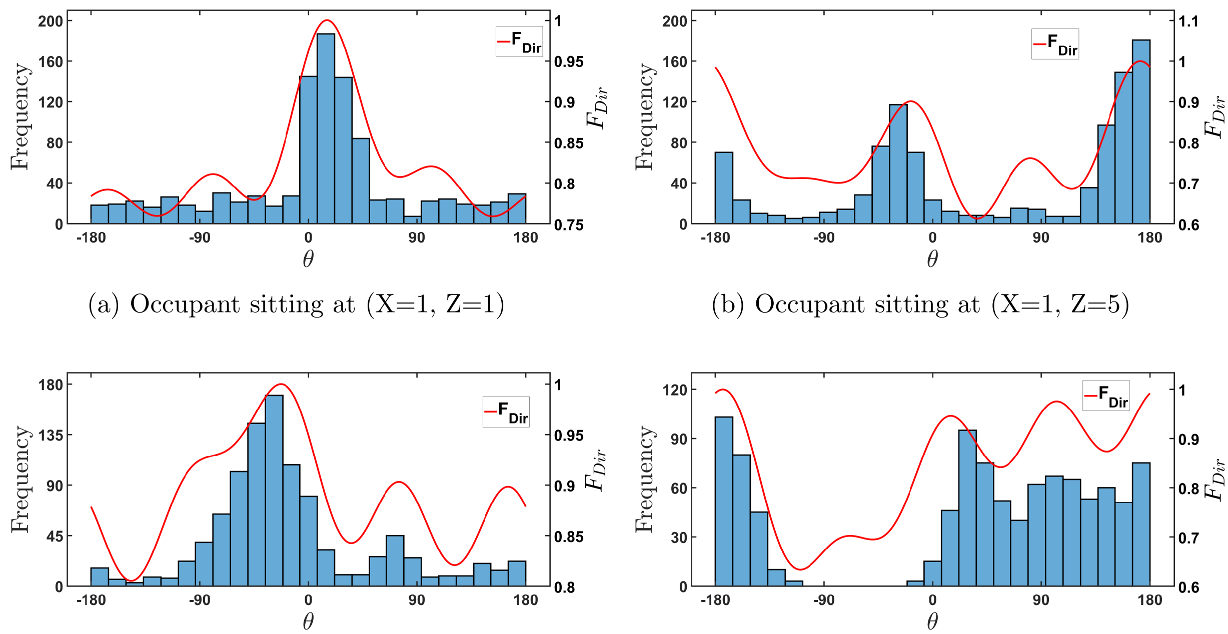

(a) Occupant sitting at (X=1, Z=1)

(b) Occupant sitting at (X=1, Z=5)

(c) Occupant sitting at (X=3, Z=1)

(d) Occupant sitting at (X=3, Z=5)

Figure 21: Histograms of the best angle of flow as predicted by the RL agent

In Figure 21(a) we have also plotted the F_{Dir} function for the occupant sitting at (X=1, Z=1). The comparison of this plot (ground truth) with the histogram of direction of airflow deemed to be optimal by the RL agent reveals that the RL agent is more likely to choose airflow directions with greater F_{Dir} value than those that have a smaller F_{Dir} value. This indicates that in this particular case on average the agent has been able to choose more energy efficient angles despite the limitations imposed on the length of the learning phase and the complexities of the environment as discussed before.

In Figure 21 similar results for three other occupant locations (as shown in Figure 20) have also been presented along with the F_{Dir} function for each particular case. As can be seen in the histogram of each case, the agent has been successful in choosing effective directions of flow more than the unfavorable ones. An interesting observation is that the agent's likelihood of choosing a direction of airflow as the best airflow direction is higher in cases where there is a more dominant airflow direction. For instance, in locations at the right-hand side of the room (Figure 21(a) and Figure 21(c)) the F_{Dir} has a more pronounced peak. In these locations, the RL control agent is more likely to choose the best direction of airflow than the cases where the best airflow direction does not have a pronounced maximum F_{Dir} value. The overall success of the RL agent in constructing an appropriate understanding of the optimal airflow direction despite the complexities and performance constraint imposed on the agent is indicative of the potential of the proposed control framework as an autonomous air-flow direction controller.

3.5 Conclusion

In this paper we have introduced the notion of enabling the HVAC system to control the direction of airflow coming into the room at the point of diffusion as apposed to the conventional uniform diffusion approach. We studied the potential energy savings that can be achieved by such expansion of control capabilities. We have further investigated the possibility of using Reinforcement Learning to enable autonomous control of the direction of airflow subject to privacy constraints. Our analysis indicates the existence of considerable potential for energy savings with the addition of the airflow direction control capacity to the HVAC control capabilities. Our observations indicate that the Reinforcement Learning algorithm could enable the HVAC system to autonomously control the direction of airflow under the constraints of privacy. Reinforcement Learning has further been able to demonstrate effective performance in terms of briefness of the learning period. In our future studies we will further investigate the potential of RL as a control algorithm in real-world experiments.

3.6 Acknowledgment

This material is based upon work supported by the Institute for Critical Technology and Applied Science (ICTAS) at Virginia Tech. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ICTAS.

4 Research Study 3: Active Diffusor, the Next Step for User-Centric HVAC Systems

Authors: Milad Abedi[1], Farrokh Jazizadeh[1]

[1]Via Department of Civil and Environmental Eng., Virginia Tech, Blacksburg, VA, 24061, USA

Key Words: HVAC Control, Directional Airflow, Thermal Comfort, User-Centric HVAC systems, Energy Consumption

4.1 Introduction

During the course of the last few decades, numerous research efforts have been dedicated to enhancing the operational efficacy of the HVAC systems in terms of their ability to improve the occupant thermal comfort while minimizing the energy consumption. Part of the incentives for researchers in this field rises from the potential for considerable energy savings by improving the performance of existing HVAC systems. According to US Energy Information Administration [1], buildings in the US are responsible for 39% of the energy consumption. The shares of the HVAC systems from the consumed energy in the commercial and residential buildings are 44% [2] and 48% [3] respectively, meaning that any improvement in the energy performance of the existing HVAC systems can lead to considerable savings in the total energy expenditure. The other part of the incentives for the aforementioned research efforts originates from the inefficiencies of the current HVAC systems in promoting occupant thermal comfort. In their survey of 215 building in US, Canada, and Finland, Huizenga et al. [57] found that only in 11% of the buildings the percentage of satisfied occupants was above 80%. Their findings further indicate that on average, the number of dissatisfied occupants is greater than the number satisfied occupants.

In order for these objectives (enhanced energy efficiency through improved Human-Centric Control Strategies) to be realized, researchers have endeavored to propose enhanced HVAC control frameworks. The level of change that these efforts have introduced into the existing HVAC control framework varies considerably from one approach to another. In one line of research, researchers have tried to conserve the existing control logic to the greatest extent possible. Under the existing control paradigm, the HVAC system will continue to condition the air in the environment until the temperature set point requirements at the location of the controlling thermostat are satisfied. However, the temperature set point of the controlling thermostat is not necessarily reflective of the thermal comfort tendencies of the occupant. Firstly, the ambient temperature at the location of the controlling thermostat is not necessarily the same as the ambient temperature at locations in the strict vicinity of the occupant. Moreover, the thermal comfort of humans has been shown to depend on a variety of parameters other than the ambient temperature, such as clothing insulation, metabolic rate, and air speed [30]. In order to compensate for these shortcomings, some researchers have proposed control frameworks, in which the temperature settings of the controlling thermostat would be driven in such a way that the thermal comfort preferences of the occupant are taken into account (e.g. [4–6, 58, 59]).

Current HVAC systems usually serve multiple rooms in a building, while only having a controlling thermostat in one of the rooms. Moreover, the actions that the HVAC systems are capable of taking have also been limited to actions that simultaneously affects multiple rooms (i.e. directing the conditioned air to all the rooms). Because of such limitations in action capabilities, the HVAC system will not have accurate knowledge of the demand in individual rooms and it will not be capable of responding to room-level demands. As a result, if one occupant in one of the rooms feels uncomfortably hot and lowers the thermostat temperature set point, the HVAC system will in turn direct the conditioned air to all of the rooms regardless of their occupancy state and preferences until the thermostat requirements have been satisfied. An intuitive solution to this problem has been to divide larger zones into a number of smaller zones where the HVAC system would be capable of directing the conditioned air into each zone independently. Once this level of control has been achieved, the system could be augmented with room-level sensors that would allow for detection of occupancy to be used in the control framework.

Notwithstanding the aforementioned advancements in the operations of the HVAC systems, there is still room for improvement in the action capabilities of the HVAC system. Under the current HVAC system design, the air in the entirety of the room will be uniformly conditioned until a stopping criterion (e.g. thermostat temperature set point) has been satisfied (Figure 22(a)). However, one might not need to condition the air in the entirety of the room to achieve occupant comfort. It might be possible to reach occupant comfort by conditioning the air only in the vicinity of the occupant by enabling control over the direction of airflow at the point of the ceiling air-vent (Figure 22(b)). In this study, we have proposed an augmentation to HVAC systems, where the control capabilities of the HVAC system will be extended to include control over the direction of airflow at the point of the ceiling air-vent. To achieve this objective, we have designed and prototyped a robotic air-vent (we will call it the active diffusor) that can enable airflow direction control. The proposed active diffusor has been designed and built with the same dimensions as the existing uniform diffusors, thereby allowing for it to be directly installed in place of the existing facilities. We have installed the active diffusor in an office room on the campus of Virginia Tech to study the effects of the active diffusor on the temperature distribution in the room, in comparison with the traditional uniform diffusors. To do so, we have installed a network of ambient temperature sensors uniformly distributed inside the within the room. By increasing the temperature inside the room to an uncomfortable 28 °C, we were able to compare the performance of the active diffusor with the traditional uniform diffusors in bringing the occupant to a state of thermal comfort.

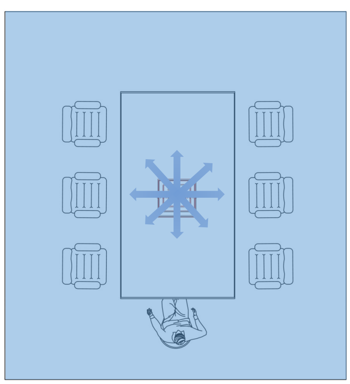

(a) Uniform Diffusion

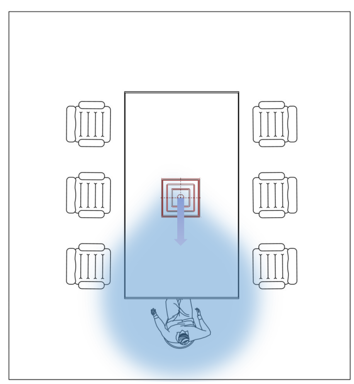

(b) Targeted Diffusion

Figure 22: The conventional and the proposed diffusion strategies

4.2 Research Background

In the recent decades, a great body of research has been dedicated to enhancement of the operational efficiency of HVAC systems in terms of energy consumption and occupant comfort. One prominent and recurring theme among these efforts is the notion of user-centric HVAC systems. In the conventional approach to HVAC control, the representation of the human side in the control framework was through the temperature set point of the controlling thermostat. Moreover, in the conventional HVAC systems, action capacities of the system were limited to zone-level on/off operations, where each zone could potentially consist of multiple spaces. In the user-centric paradigm, the underlying operational logic of the HVAC system has been altered by directly including the human-comfort and occupancy in HVAC system operations.

In order for a human-in-the-loop HVAC framework to be realized, the HVAC system must be capable of constructing a detailed understanding of the human-side demand. In determination of demand, the HVAC should be able to detect the state of occupancy as an essential step. Researchers have employed a variety of sensing technologies for occupancy detection including RFID based systems [60], WiFi based frameworks [61], and vibration based approaches [62]. After determination of occupancy, the user-centric HVAC system must be able to understand occupant thermal comfort. Since its initial introduction, the Predicted Mean Vote (PMV) model [30] has gained a considerable amount of attention from the research community and has been used in numerous studies as a means of understanding occupant thermal comfort tendencies [63–65]. Notwithstanding its widespread popularity, by design, the PMV model is reflective of the thermal comfort tendencies of a group of occupants, and as such it does not necessarily reflect the thermal comfort tendencies of the specific occupants [66, 67]. In efforts to address the inaccuracies that arise from such generalizations, researchers have proposed personalized frameworks to help capture the individualized thermal comfort attitudes. To that end, numerous studies [13, 15, 17, 68, 69] have proposed user-feedback based frameworks where the occupants directly inform the HVAC system of their thermal comfort state. Researchers have also studied other, less burdensome approaches for evaluation of the thermal comfort state of the occupants through measurement of such physiological parameters as subcutaneous blood flow [12, 37, 40, 70], pulmonary activity [38, 39], skin temperature [41, 71–73], and heart rate [74].

While having an accurate knowledge of the occupants' thermal comfort state is a necessary piece of information for the realization of a user-centric HVAC system, it is not sufficient. To be truly user-centric, the HVAC system must also be capable of taking appropriate actions in accordance with observations regarding occupant thermal comfort. In the conventional HVAC control framework, the HVAC system directs the conditioned air into the environment so as to satisfy the temperature requirements of the controlling thermostat. In efforts to render the existing HVAC systems user-centric without major alterations in the operational design of the HVAC systems, Al-Sanea and Zedan [75] have aimed to optimize the fixed monthly thermostat set point of the HVAC system while accounting for occupant comfort. Other researchers [4, 5, 76, 77] have also proposed comfort-based control frame works for driving thermostat temperature set points.

The improvements brought on by these frameworks notwithstanding, there are still unaddressed limitations present in these approaches. In many HVAC systems, one thermostat and one actuator serve each zone which often consists of multiple spaces. Moreover, the response of the HVAC system to the requirements of a thermostat is to direct the airflow to all of the rooms in that zone. As a result of this setup, the differences in temperature in different rooms of the zone, the variations in the thermal comfort preferences of the occupants in different rooms, and the occupancy states of the rooms are not being taken into account. To help address these shortcomings, Lin et al. [78] have endeavored to enhance the control of the existing single actuator per zone systems by using data acquired from multiple sensors (one per room) without any infringement on the original action capabilities of the HVAC system. Other researchers [52, 79, 80], in a departure from the practice of minimal-intrusion into the original design of the HVAC mechanics, have studied the potential benefits of a higher control resolution by dividing larger zones with multiple rooms into a number of smaller zones with on/off capabilities for each zone, thereby permitting for incorporation of space occupancy state in HVAC control Xu and Wang [81]. There are also numerous commercial products (*smart vents*) that can replace the existing vents in residential buildings to facilitate binary control capabilities [82, 83].

The studies discussed so far have operated under the assumption of a uniform diffusion setting. In the uniform diffusion paradigm, the incoming airflow into the room will be uniformly distributed in all directions at the point of the ceiling air-vent by means of conventional uniform diffusors. By forgoing the assumption of uniformity, in an effort to utilize the non-uniform nature of the temperature distribution to achieve higher levels of thermal comfort, Lo and Novoselac [84] have used localized airflow to create areas of comfort within the room in an energy-efficient manner. In another study, Wang et al. [85] have aimed to take into account the spatial dynamics of the HVAC system operations by tracking dynamic spatial occupancy distribution within a room so as to control the HVAC system accordingly.

By casting aside the assumption of a uniform environment, in this study we have proposed further modifications in the existing HVAC systems to enable control over the direction of airflow at the point of the ceiling air-vent. In our previous simulation study [32] we have demonstrated the existence of a considerable potential for energy savings by enabling control over the direction of the airflow at the point of the ceiling air-vent. In this study, by enabling the HVAC system to control the direction of airflow in a real-world experimental setting, we have sought to study the feasibility, thermophysical impact, and the potential benefits of creating a local comfort climate surrounding the occupant by adjusting the direction of airflow. To do so, we have designed, prototyped, and implemented a the active diffusor which allows for control over the direction of airflow at a ceiling air-vent in office buildings.

4.3 Methodology

4.3.1 Directional airflow

4.3.1.1 Concept In the conventional HVAC systems, the diffusors, as the name implied, are not intended to be an active part of the control framework. They have been designed to distribute the air uniformly in a given space for thermal and ventilation conditioning. However, with advancements in the field of user-centric HVAC control and the increased granularity of the control process, the idea of diffusors as mere stationary pieces could be rethought. Currently, multiple commercial products under the common description of *smart vents* [82, 83] are present in the market. These devices have been designed to replace the existing diffusors in residential buildings to add such capabilities as on/off functions. These added capabilities, coupled with technologies that allow for detection and prediction of occupancy have facilitated higher levels of user-centricity in the HVAC control process.

In this study, we have sought to further augment the control mechanism of HVAC systems to create additional adaptive control capacity over the direction of airflow, in addition to the binary on/off functionality. The resulting extensions in the control capacity could be utilized to create targeted areas thermal comfort surrounding the occupant as opposed to the conventional practice of diffusing the conditioned air uniformly into the space until the entirety of the room is thermally comfortable. The obviation of the need to condition the entirety of the space, could result in savings in energy consumption.

To fully utilize the added control capabilities towards targeted response and energy efficiency, the HVAC control framework must also undergo certain modifications. For instance, depending on the distance of the occupant from the active air-vent (capable of creating directional airflow) the energy efficient values of the supplied air speed and temperature could vary. As such, in an ideal condition, the HVAC system is capable of dynamically changing these variables in accordance with the preferences and the location of the occupant.

Additionally, in the conventional uniform diffusion approach, from all the factors that would affect the thermal comfort state of the human subject, the HVAC systems are only able to control the temperature. However, with the addition of directional airflow capabilities, the synergistic combination of the expanded control over the direction of airflow and the ability of the HVAC system to adjust the airflow rate would result in a capability to also control the airflow speed at the strict vicinity of the occupant. Given that the air speed is one of the influential factors in determining the thermal comfort state of the occupants [30], the ability of the system to influence it could result in further savings in a user-centric control framework.

4.3.1.2 The Envisioned Framework In order for the potential benefits of directional airflow to be realized, the HVAC system must also be capable of effectively controlling the direction of airflow. While it would be possible to query the occupant about the best direction of the airflow, such an intrusive approach would neither be operationally feasible, nor effective as the occupant might not necessarily be aware of the best airflow direction. As such, autonomous control of the airflow direction by the HVAC system is a necessity for the realization of the potential benefits of directional airflow.

In our previous study [32], we observed that the location of the occupant inside the room is an invaluable piece of information when it comes to controlling the direction of airflow. As such, one could utilize indoor localization technologies that allow for knowledge of occupant location. However, the utilization of indoor localization technology is associated with a number of privacy concerns, as well as additional sensing infrastructure.

Alternatively, certain control approaches could be used that would help obviate the need for information regarding the location of the occupant. One possible candidate for such a control framework is the utilization of Reinforcement Learning (RL) algorithm which makes it possible to control the direction of airflow based on a reward given to the control framework for its actions. In the RL setting, the reward is indicative of how favorable the action in question was such that more favorable actions would lead to greater rewards and vice versa. By defining the reward to be a function of occupant thermal comfort and energy consumption the RL algorithm could be used to achieve these objective in controlling the direction of airflow. The authors are currently studying the feasibility of utilizing such an autonomous control framework for control of the airflow direction in another paper.

For the purposes of this study, we have assumed that only one occupant will be present in the room at any moment in time. However, The directional airflow could also benefit rooms with multiple occupants, since the direction of airflow could be dynamically controlled to serve multiple occupants in the room. Another consideration for the envisioned framework is the ventilation requirements. Under the uniform diffusion, the fresh air is diffused through the room to ensure that the air quality throughout the room is at an acceptable level. However, with directional airflow and the quality of the air at area that are far from the targeted location could deteriorate. This concern could be addressed by periodically directing the airflow to those locations.

4.3.2 Active Diffusor Design and Prototyping

In this study, we have sought to build and test a devise that enables us to investigate the potential benefits of directional airflow in terms of improving occupant thermal comfort and HVAC energy efficiency, as discussed in the previous sections. To achieve this, we embarked on a process of designing a device (active diffusor), that would make it possible for the HVAC system to control the direction of airflow. In the process of building the active diffusor we came up with a number of potential designs.

Presented in Figure 23(a) is the first prototype of the active diffusor we designed. In this design, flexible air ducts, the orientation of which could be changed to direct the airflow at the direction of interest were utilized. To adjust the orientation of the tubes, linear actuators were used to move the upper frame of the devise, to which the tubes have been connected. The lower part of the tubes are connected to gyroscope bearings (Figure 23(b)), making it possible for the orientation of the tubes to be freely adjusted. We proceeded with building the first prototype, and during the building process we realized that this design could be improved in a number of ways. To begin with, the tubes needed to be made out of material that was flexible enough to be easily bent in the direction of interest, yet strong enough to hold the circular tube cross-section shape without collapsing. A number of available products in the market were experimented with and the conclusion was reached that the tested materials were not suitable for this application. Moreover, because of the relatively high volume of steel rails used in this design, the weight of the final product was also a concern, specially considering that the device was intended to be mounted on the existing suspension ceilings which have not been designed to carry heavy loads. After facing these challenges in building the first prototype, we proceeded with the design of a second prototype to address the concerns that were raised during our first attempt. Our design of the second prototype was driven by a number of criteria imposed to ensure the feasibility of the active diffusor as an economically and functionally viable product.

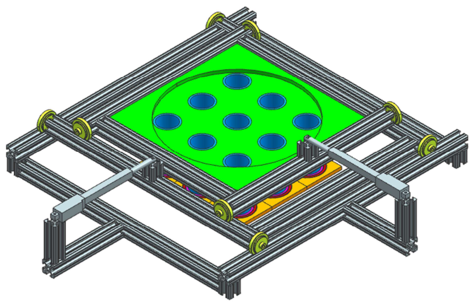

(a) Overview of the first prototype of the active diffusor

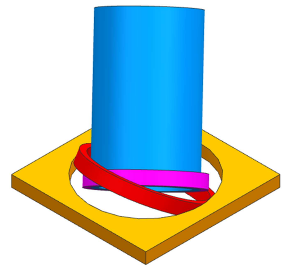

(b) Gyroscope bearing mechanism

Figure 23: The first prototype of the active diffusor

The first and most important criterion in the design of the active diffusor was its capability to replace existing uniform diffusors and direct the airflow at the desired direction. To that end we decided to use louvre blades as a means of directing the air flow. Louvre blades have been used for years as a means of adjusting the direction of airflow in the car industry and in the manually adjustable air-vents for HVAC systems. In Figure 24(a) a side view and in Figure 24(b) a top view of the two layers of louvre blades as well as the connector pieces have been displayed. In order for the active diffusor to be installable in place of the existing uniform diffusors, we have chosen the outer dimensions of the active diffusor to be $2ft \times 2ft$ $(0.61m \times 0.61m)$, which is the typical grid size of suspension ceiling in the US. Therefore, one of the questions being addressed in this study is with regards to the feasibility of the proposed design in effectively directing the airflow jet towards the targeted location.

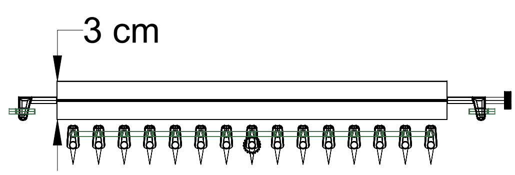

(a) Side view

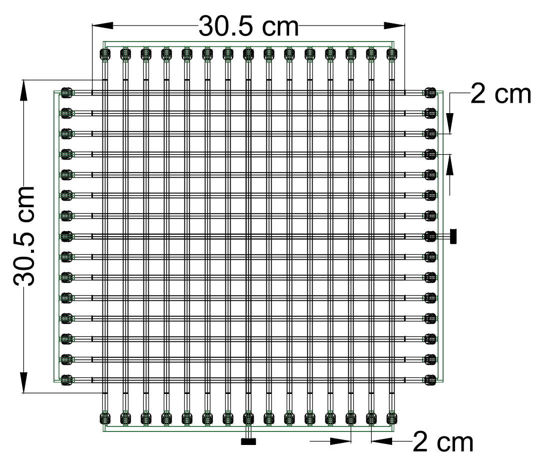

(b) Top view

Figure 24: Side view of louvre blades and connector pieces

Another design criterion that drove our design was the cost considerations. In the current design, we will need one actuator each one of the two louvre blades. If it were not for the cost considerations, instead of two layers of connected louvre blades we could have divided the opening into 4 equal sections with 4 sets of double-layered louvre blades operating independently of one another, with an increased number of actuators. While such a design would result in further control granularity (ability to simultaneously target multiple locations in the room), the added complexity would result in higher costs.

For actuation, we have chosen to use two stepper motors, one for each layer of louvre blades. An image of the active diffusor has been presented in Figure 25. As can be seen in Figure 25, the stepper motors have been used in conjunction with worm-gears to ensure the blades would retain their position. To enable remote control over the actuators, we have used a Raspberry-Pi computer, capable of connecting to the existing wireless network. In order to power this setup, we will need a 12-volt power supply for the actuators and a 5 volt power supply for the Raspberry-Pi computer. Given the presence of wiring for lights in the ceiling, powering the active diffusor does not require excessive extra wiring.

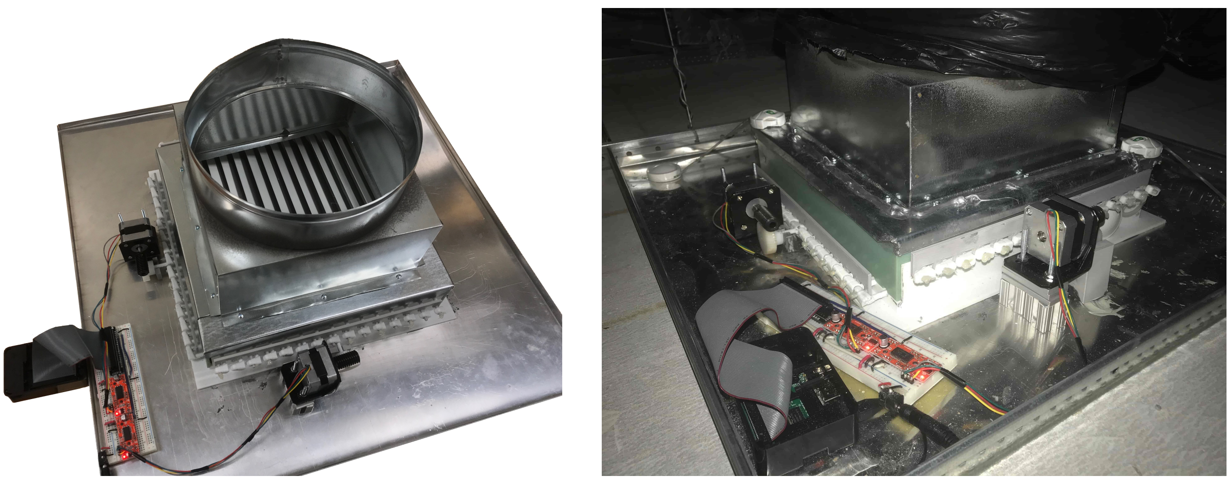

(a) Before installation

(b) After installation

Figure 25: The active diffusor

Following the installation of the active diffusor in the room environment, we sought answers to a number of research questions through experimentation. The first question was whether the implemented design for the active diffusor could successfully target the occupants seated in the room. To answer this question, we have studied the impact of the active diffusor in the change of temperature at the targeted occupant location compared to other non-targeted locations. The second question to be answered, is whether the targeted airflow approach will lead to savings in terms of energy consumption. In order to answer this question, we have conducted experiments using the conventional uniform diffusors so as to produce a baseline against which the efficacy of the targeted diffusion approach is to be compared. In this study we have adopted the duration of flow required to bring the occupant to a state of comfort as the comparison metric.

4.3.3 Experimental setup

In order to evaluate the performance of the active diffusor, we have installed it in place of an existing uniform diffusor in a computer room located at Virginia Tech with $7.5m \times 4m$ dimensions. We have also installed a network of ambient air temperature sensors, uniformly distributed through the room to capture the temperature distribution created by the active diffusor. In Figure 26, the plan of the office room, as well as the distribution of the sensors can be seen.

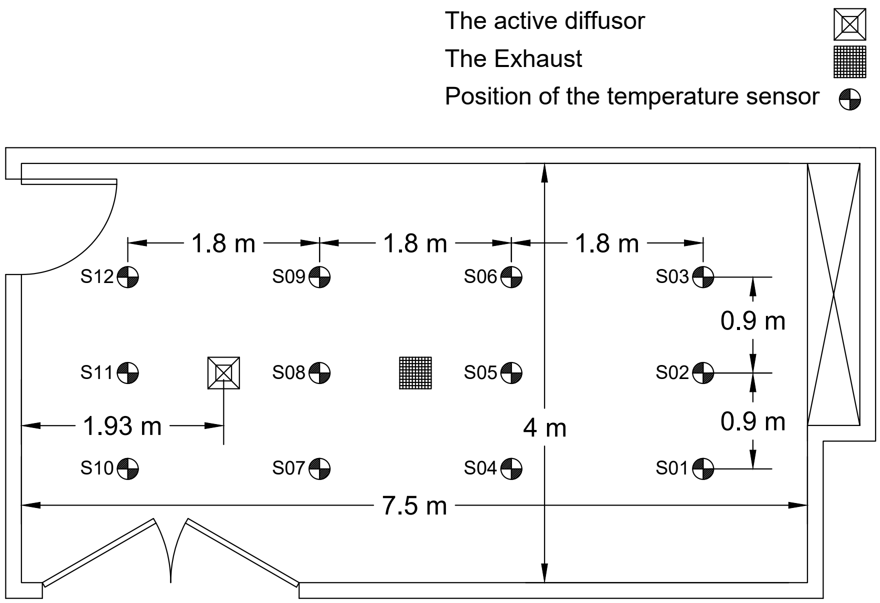

Figure 26: The plan of the test room and the distribution of sensors

As shown in Figure 26, a total of 12 temperature/humidity sensors have been installed in a 3×4 grid at a height of 80 cm from the floor, corresponding to the approximate average height of the abdomen area of the occupants seated in an office chair. The sensors used in this experiment were digital DHT22 temperature and humidity sensors, although in this setup we only used the temperature readings. For data acquisition we have used the Arduino Mega 2560 Rev3 micro-controller board. In Figure 27 an image of the experimental setup has been shown.

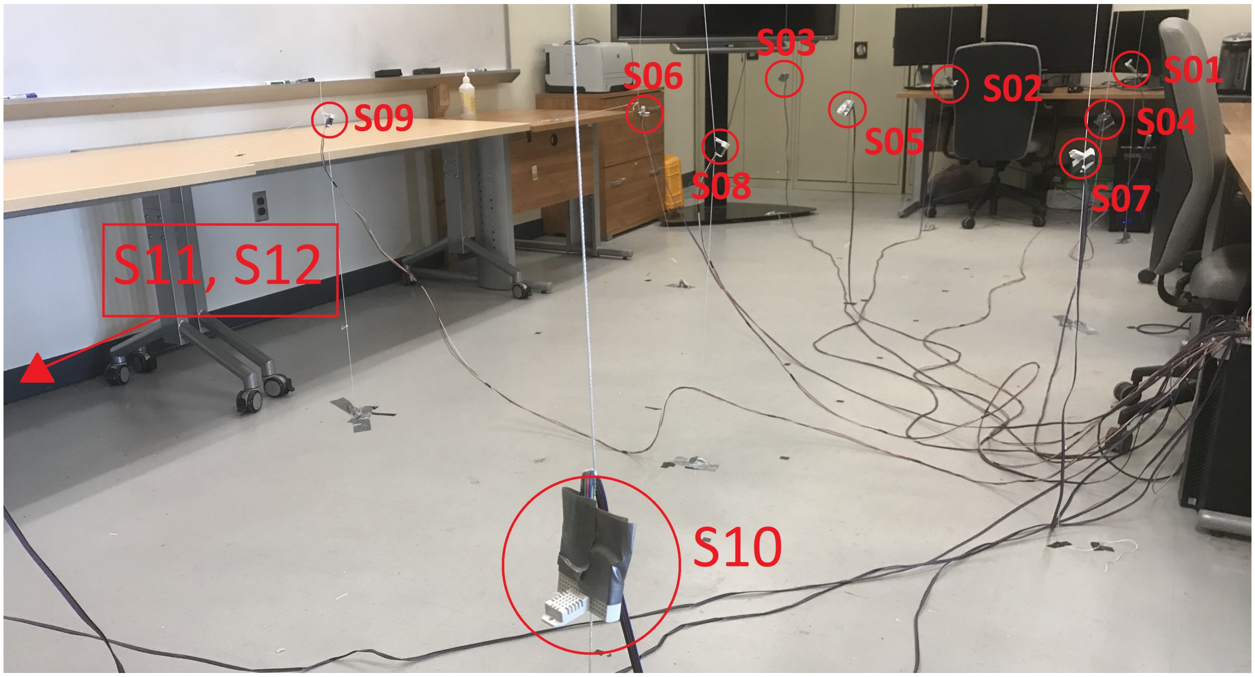

Figure 27: Image of the sensor setup

In order to study the ability of the active diffusor in creating a local comfort zone surrounding the occupant (represented by temperature sensors), we used 4 heaters and one fan to increase the temperature of the room to a uniform 28 °C. Then we create a local comfort zone for the imaginary occupant sitting at the location of each temperature sensor by directing the airflow towards that point via the active diffusor. The rate of temperature drop at the targeted position, compared to other non-targeted positions, will be observed to determine the efficacy of the active diffusor in creating targeted flow. We performed this experiment for each of the 12 temperature sensors, each representing a possible location for the occupant. The uniform distribution of these sensors through the room helps obviate any assumptions regarding the location of the occupant and any possible biases that could arise from such assumptions. In order to develop a baseline against which to compare the performance of the active diffusor, we have also evaluated the performance of the conventional uniform diffusors in a similar experimental setting. The duration of time required to bring the occupant to the state of thermal comfort in the case of active diffusion and uniform diffusion will be studied and compared so as to serve as an indicator of energy saving potential of the active diffusor. Note that in all of our experiments we have not controlled the inner functions of the HVAC system, and the HVAC system has relied on its own thermostat based control logic.

In this paper, for evaluation of the thermal comfort state of the occupants, we have used a thermal comfort model offered by Daum et al. [18] which relates the ambient temperature at the vicinity of the occupant to the probability of being thermally comfortable (Figure 28). According to the ASHRAE 55 standard [7], the lowest acceptable percentage of satisfied occupants (probability of being comfortable) is 80%. By adopting the same minimums of a thermally comfortable environment, we have defined the state of comfort to be a state where the occupant has an 80% chance of being thermally comfortable (equal to or above $TC_1 = 0.8$

as shown in Figure 28). Starting from an initial ambient temperature of $T_0 = 28\,°\mathrm{C}$, in Figure 28 the initial value of the thermal comfort index (TC_0), the minimum changes required in thermal comfort index of the occupant (ΔTC) and in the ambient air temperature (ΔT), as well as the corresponding final ambient air temperature (T_1) have been presented.

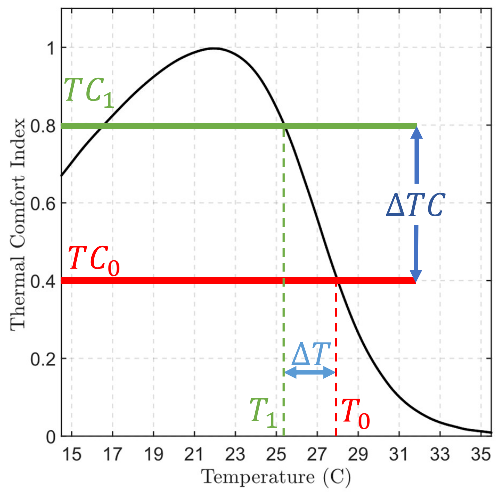

Figure 28: The model of thermal comfort offered by Daum et al. [18]

4.4 Results

Prior to conducting our experiments, we tested the data acquisition system to ensure a consistent and sufficient flow of data acquisition capabilities. For our experiments, we were able to read the temperature every 3.805 seconds (a constraint imposed by the data acquisition setup). The sampling period (time elapsed between two samples) was very consistent, with minimum and maximum of 3.796 ans 3.811 seconds respectively. We determined that this rate of sampling was sufficient for our application and there was no need to use more advanced data acquisition hardware to obtain higher sampling frequency. While we expected a normal distribution for the sampling period, our observations indicated a bi-normal distribution for these values (Figure 29). Given that the all of the values for the sampling period are very close to the average value, we accepted the non-normal distribution of the sampling period as is and did not look further into the issue.

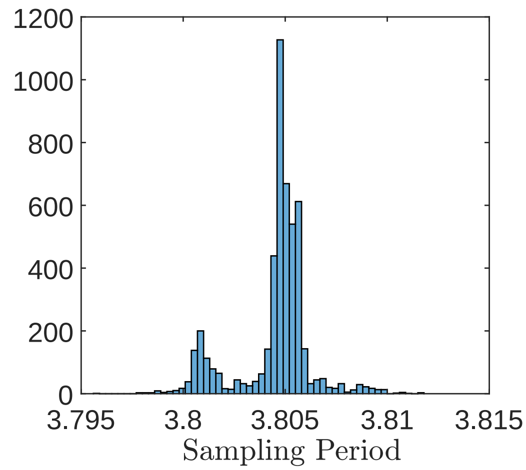

Figure 29: Histogram of the sampling period (time elapsed between two samples) for sensor readings

In our first experiment, we have studied the performance of the conventional uniform diffusors in creating a comfortable climate in the vicinity of the occupant (represented by temperature sensors). To do so, the temperature of the room was risen to an uncomfortable (according to Figure 28) 28 °C . Then the HVAC system started to condition the air in the room while diffusing the conditioned air using a conventional uniform diffusor which results in a sharp angle between the airflow jet and the ceiling. The temperature readings of each of the 12 sensors shown in Figure 26 have been plotted in Figure 30. As can be seen in Figure 30, the trends of the change in the temperature are consistent for different occupant locations.

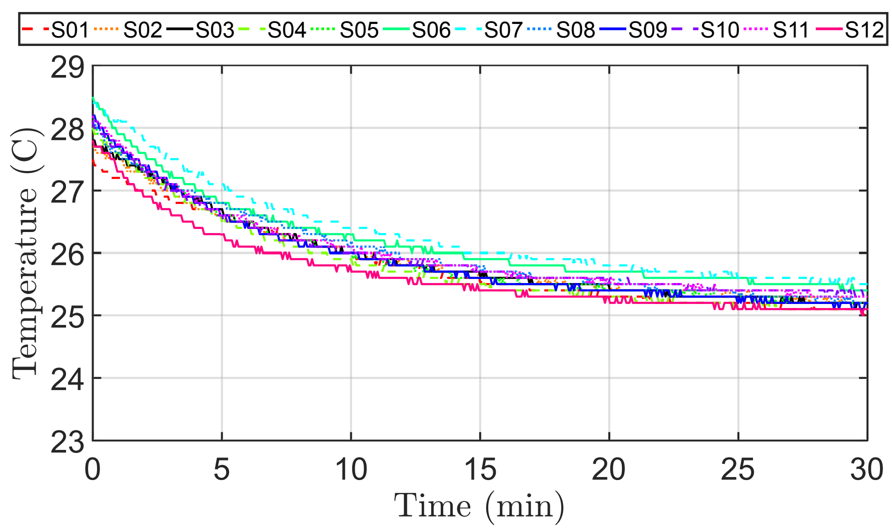

Figure 30: Ambient temperature under uniform diffusion at the location of temperature sensors

For the next set of experiments we replaced the existing uniform diffusor with the active diffusor. As was the case for the previous experiment the initial temperature of the room was set to an approximately uniform temperature of 28 °C. In order to evaluate the ability of the active diffusor in creating local comfort zones surrounding the imaginary occupant(represented by a temperature sensor), the direction of the airflow was adjusted towards the corresponding sensor positions. This experiment was performed for each one of the 12 occupant sitting positions. Changes in the ambient temperature through the room, when sensor positions S01 through S06 have been targeted are presented in Figure 31.

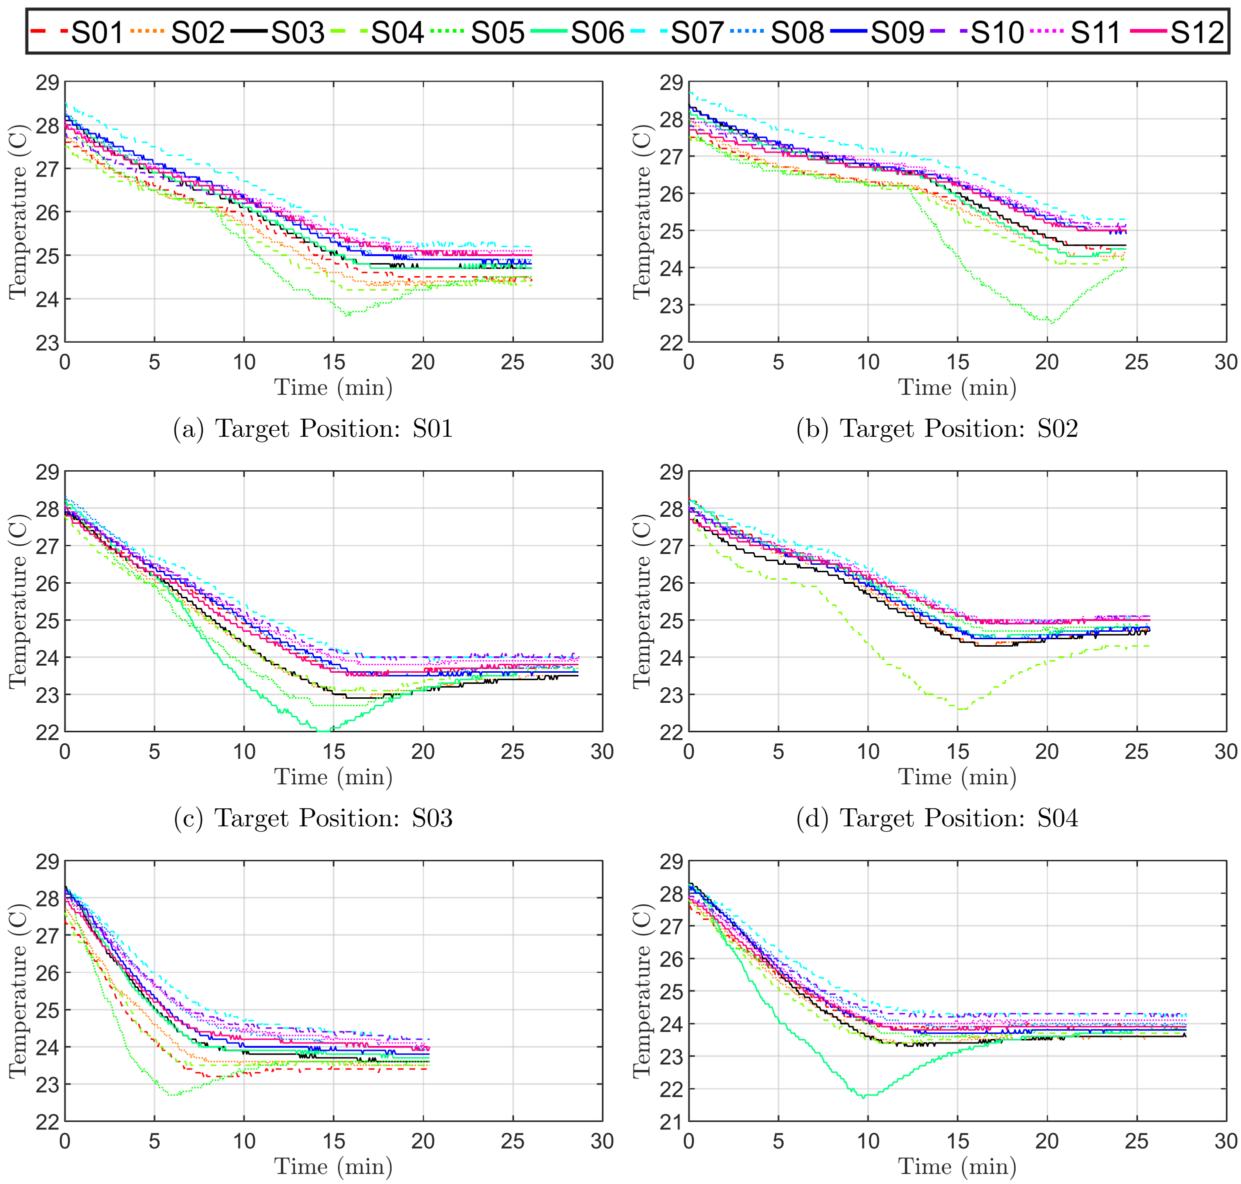

(a) Target Position: S01

(b) Target Position: S02

(c) Target Position: S03

(d) Target Position: S04

(e) Target Position: S05

(f) Target Position: S06

Figure 31: Targeted airflow experiment for S01-S06 sensor positions

As shown in Figures 31(d), 31(e), and 31(f), when the active diffuor aims to create a local comfort zone in these positions by directing the airflow towards them, the change in temperature at these points happens faster than non-targeted positions. However, as can be seen in Figures 31(a), 31(b), and 31(c), while the active diffusor has aimed to target sensor positions S01, S02, and S03, the change in temperature at these locations was slower than the change of ambient temperature at non-targeted positions of S04, S05, or S06. This is due to the fact that S01, S02, and S03 are at a farther distance from the active diffusor and when the active diffusor aims to target these positions, it also ends up targeting S04, S05, and S06 due to the dispersion of the airflow jet. Given the greater relative distance between sensor positions S01-S03 and the diffusor, when targeting these positions, due to the greater dispersion caused by the greater travel distance of the airflow jet, other sensor positions also get affected. As such, in these cases, the difference in temperature drop for the targeted sensor position compared to other non-targeted sensor positions is not as significant.

By contrast, when targeting sensor positions that are closer to the active diffusor (i.e. S07-S12) the difference between the temperature drop at the point of the targeted sensor versus other non-targeted sensor positions is more pronounced as shown in Figure 32. This is due to the lesser dispersion of the airflow jet due to the shorter travel distance to the targeted position.

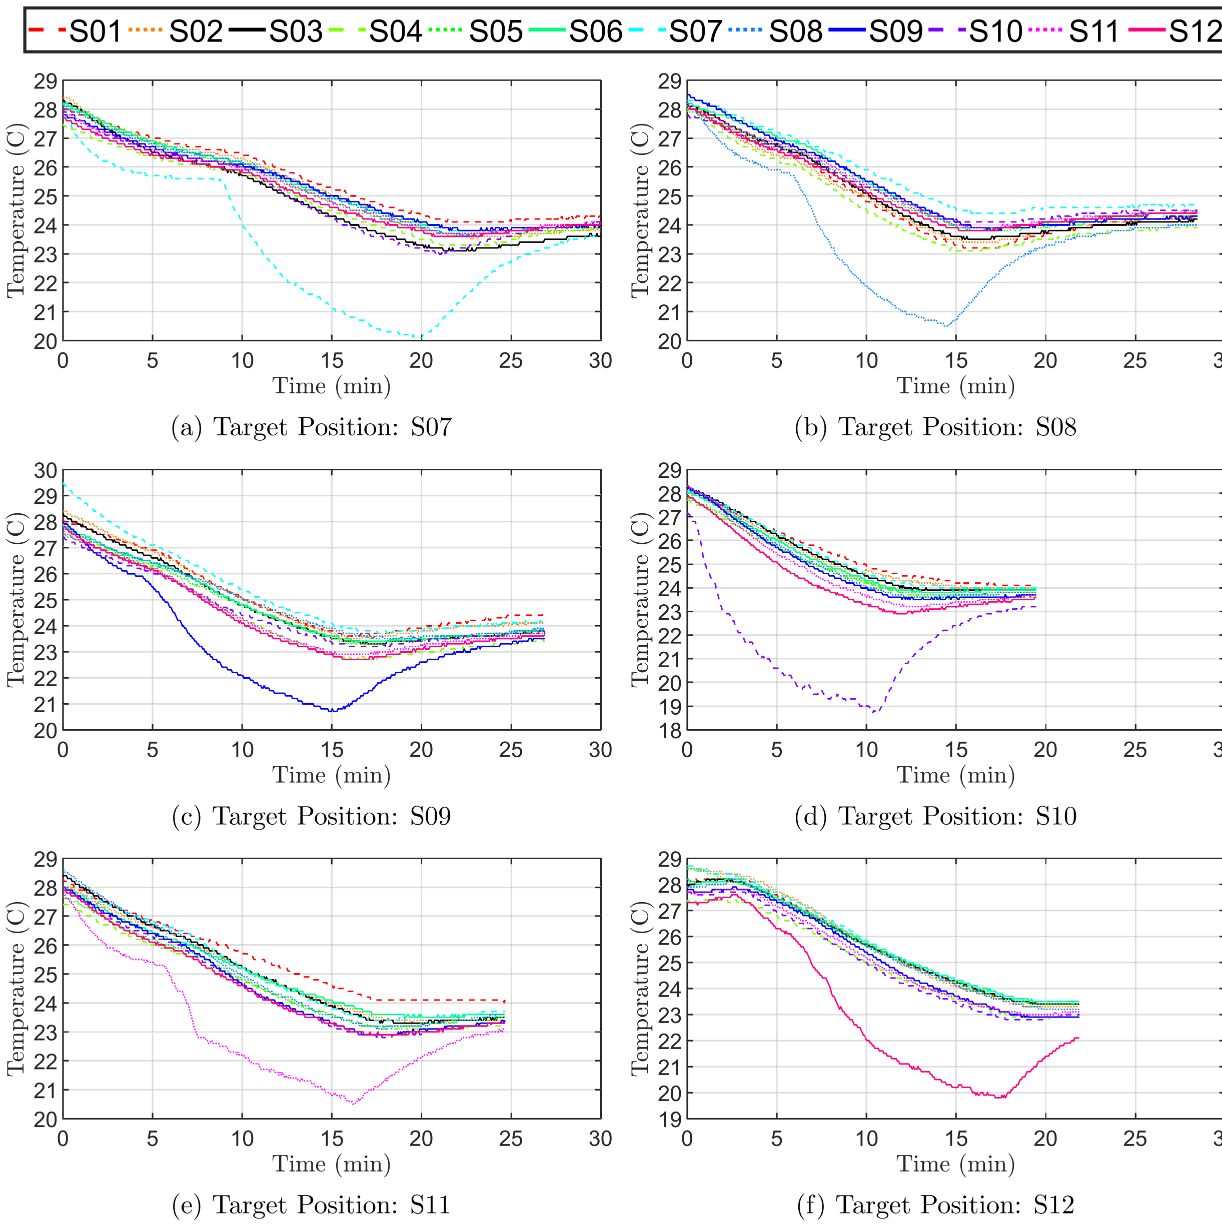

(a) Target Position: S07

(b) Target Position: S08

(c) Target Position: S09

(d) Target Position: S10

(e) Target Position: S11

(f) Target Position: S12

Figure 32: Targeted airflow experiment for S07-S12 sensor positions

One of the contributions of this study is the presentation of a design for the active diffusor, that is capable of effectively targeting the airflow jet towards the direction of interest. The stark difference in the rate of change in ambient temperature between the targeted and non-targeted positions where the targeted position experiences a faster rate of temperature drop in all experiments except for experiments where positions S01-S03 were targeted, is indicative of the ability of the prototyped active diffusor to effectively direct the airflow jet at the direction of interest so as to create a locally comfortable zone surrounding the occupant.

As mentioned in the previous section, another objective of this study was to demonstrate that directional airflow is a more energy efficient approach when compared to the conventional uniform diffusion scenarios. In their 2014 paper, Ghahramani et al. [4] demonstrated that a linear relationship exists between the electricity consumption of the HVAC system and the airflow volume. Thus, in order to investigate the potential for energy savings, we studied the duration of time required to bring the imaginary occupant to a state of comfort for both the uniform diffusion and targeted diffusion scenarios, when the initial temperature in the room is an approximately uniform 28 °C. In Table 3, the required durations of airflow until comfort for each occupant location (i.e. sensor position) has been presented for both uniform and targeted airflow. Moreover, we have also presented the percentage of reduction in the duration of airflow that can be achieved via targeted airflow compared to uniform airflow. An average reduction of 64% was achieved in the duration of flow required to reach a state of thermal comfort (time-to-comfort). As such, the reductions in the duration of airflow achieved by active diffusion as shown in Table 3 can be taken to mean a proportional reduction in the energy consumption. Moreover, the ability of the active diffusor to create a state of comfort in a shorter period of time also means that the occupant will be uncomfortable for a shorter duration of time, thereby enhancing the overall thermal comfort experience.However, it is the opinion of the authors, that further studies are required to better understand the effect of directional airflow on the thermal comfort state of the occupants.

Table 3: Duration of airflow (minutes) required to reach the state of thermal comfort for uniform and targeted airflow scenarios, and the percentage of reduction achieved in duration of airflow by enabling directional airflow

	Sensor Position											
	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	S11	S12
Uniform	16.9	19.5	20.0	15.3	18.5	29.0	30.8	21.4	17.9	23.0	23.0	14.4
Targeted	12.4	15.7	7.2	8	2.2	3.3	8.9	6.2	5.1	0.9	4.6	6.6
% Reduction	%27	%19	%64	%48	%88	%89	%71	%71	%72	%96	%80	%54

One interesting and unexpected observation through these experiments was the fact that even in the absence of an effective targeting framework, the active diffusor is able to reach considerable savings compared to the uniform diffusion. This is due to the fact that in the uniform diffusion approach, the conditioned air is diffused at a very sharp angle from the ceiling, conditioning the room space from top to bottom. However, regardless of the precision of the targeting functionality, in the active diffusion the airflow jet is directed towards the floor, thereby conditioning the room environment from bottom to top.

To study the energy saving potential of the active diffusor with an ineffective targeting functionality, we have studied scenarios where the distance between the intended target location and the actual target point is at least 3 meters on the floor plan. For instance, if the intended target was sensor position S01, we will use the data from experiments where sensor positions S07-S12 were targeted (these positions are at least 3m away from the intended position S01). To study the effect of the aforementioned phenomenon on the thermal environment of the room, in table 4 we have presented the average duration of airflow required for a given target sensor position to reach a comfortable temperature when the targeting functionality is ineffective (i.e. a 3m error in target localization).

Table 4: Duration of airflow (minutes) required to reach the state of thermal comfort when the targeted sensor is at least at a distance of 3m, and the resulting percentage of reduction in duration of airflow compared to uniform diffusion

	Intended Target Position											
	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	S11	S12
Distant Target	10.3	9.5	9.2	7.3	8.2	8.7	15.9	14.1	14	11.5	11.6	10.8
% Reduction	%39	%51	%54	%52	%56	%70	%48	%34	%22	%50	%50	%25

Also shown in Table 4 is the percentage of reduction in the time-to-comfort (comfortable temperature) for each intended target position. The average of percentage of reduction in time-to-comfort for this scenario is 46% compared to the uniform diffusion approach. This is indicative of the potential for energy savings compared to uniform diffusion despite the possible inaccuracies in the targeting capabilities of the framework.

4.5 Limitations and future work

In this paper we have evaluated the potential benefits of directional airflow only for one room with specific dimensions and HVAC design. It is the authors' opinion, that to better understand the broader impact of such a modification in the operational capabilities of the HVAC system, more studies in a variety of different thermal environments are required.

Another limitation of this study is the assumption that the thermal comfort state of the occupant can be accurately predicted by measuring the ambient air temperature in their vicinity. Specifically, in this study we have not studied the effect of air speed on the thermal comfort state of the occupant. In future studies, the impact of targeted diffusion on local air speed, and the effect on occupant thermal comfort should be investigated.

As mentioned in the Methodology section, we were not able to control the inner components of the HVAC systems in our experiments, and as such we had to rely on the existing thermostat-based control logic of the HVAC system. However, future studies should investigate the potential for enhanced control of the inner components of the HVAC system to better utilize the added control capacity. Specifically, the ability to control the supply airflow speed is of particular interest. With control over the incoming airflow speed, the authors could account for and control the air speed at the strict vicinity of the occupant which is an important factor in determining the thermal comfort state of the occupant.

The focus of this study has been the introduction of the idea of directional airflow and active diffusion and an illustration of its potential to improve occupant comfort and HVAC energy efficiency. However, a crucial requirement for the active diffusor is the ability of the system to autonomously control the airflow direction. In future studies the feasibility of an autonomous control framework will be studied.

4.6 Conclusion

In the present paper, a novel expansion in the action capabilities of the HVAC systems was proposed. With the envisioned modifications, the HVAC system will be able to control the direction of airflow at the point of the ceiling air-vent as opposed to the conventional uniform diffusion approach. To enable control over the direction of airflow for the existing HVAC systems, we designed and built a device called the "active diffusor" which could replace the existing uniform diffusors. We have installed the active diffusor in an office room at Virginia Tech along with a network of temperature sensors (so as to represent the occupant) to study the potential of the directional airflow in enhancing the overall occupant thermal comfort and the energy efficiency of the HVAC system. Our experiments indicate that the proposed prototype of the active diffusor can effectively direct the airflow jet towards the direction of interest. In our experiments the active diffusor reached an average of %64 reduction per cycle in the duration of airflow required to bring the occupant to an acceptable state of thermal comfort compared to the conventional uniform diffusors, translating into a considerable potential for energy savings. These results are indicative of a considerable potential for enhancing the existing HVAC systems by realization of the idea of directional airflow.

4.7 Acknowledgment

This material is based upon work supported by the Institute for Critical Technology and Applied Science (ICTAS) at Virginia Tech. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ICTAS.

5 Conclusion

In this thesis at hand, the idea of directional airflow for HVAC systems has been proposed as an improvement on the existing uniform diffusors. While in the conventional uniform diffusion approach the distribution of the temperature inside the room was assumed to be uniform, in the proposed directional airflow approach the non-uniformity of the thermal environment has been recognized and leveraged towards further savings in energy expenditure. In the proposed directional airflow approach, instead of conditioning the air in the entirety of the environment as is done in the case of conventional uniform diffusors, the HVAC system would create a local area of thermal comfort surrounding the occupant. In the first research study presented in this thesis, the potential for improvements in the energy efficiency and the overall occupant thermal comfort was investigated in a simulation setting. Per conclusions of the first research study, the directional airflow demonstrated a 59% reduction in the energy consumption per cycle, as well as improvements in the overall thermal comfort experience of the occupant via a shortening in the duration of discomfort. In the second research study, the feasibility of autonomous control of the direction of airflow by means of Reinforcement Learning was investigated. In developing the control framework, privacy considerations regarding the utilization of indoor localization technology were taken into account by operating in the absence of knowledge regarding the location of the occupant. The second research study concluded that the Reinforcement Learning algorithm could be utilized to provide an autonomous control algorithm for control of the direction of the airflow. Given the encouraging findings in the first two research study, in the third research study the functional feasibility of directional airflow as well as its potential benefits were investigated by designing and building a robotic device (called the active diffusor) that rendered the existing HVAC systems capable of controlling the direction of airflow at the point of the ceiling air-vents. The third research study concluded that the proposed active diffusor was able to effectively direct the airflow jet at the direction of interest. Moreover, the study demonstrated the capability of the active diffusor in improving the energy efficiency of the HVAC systems by 64% per cycle as well as the overall occupant comfort.

References

[1] US Energy Information Administration. (2018) How much energy is consumed in u.s. residential and commercial buildings? [Online]. Available: https://www.eia.gov/tools/faqs/faq.php?id=86&t=1

[2] U. E. I. Administration. (2016) 2012 commercial buildings energy consumption survey: Energy usage summary. [Online]. Available: https://www.eia.gov/consumption/commercial/reports/2012/energyusage/

[3] U. D. of Energy. (2014) Principles of heating and cooling. [Online]. Available: https://www.energy.gov/heating-cooling

[4] A. Ghahramani, F. Jazizadeh, and B. Becerik-Gerber, "A knowledge based approach for selecting energy-aware and comfort-driven hvac temperature set points," *Energy and Buildings*, vol. 85, pp. 536 – 548, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778814007932

[5] N. Nabil and M. Samir, "A cost-effective operating strategy to reduce energy consumption in a hvac system," *International Journal of Energy Research*, vol. 32, no. 6, pp. 543–558, 2007. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/er.1364

[6] P. Bermejo, L. Redondo, L. de la Ossa, D. Rodríguez, J. Flores, C. Urea, J. A. Gámez, and J. M. Puerta, "Design and simulation of a thermal comfort adaptive system based on fuzzy logic and on-line learning," *Energy and Buildings*, vol. 49, pp. 367 – 379, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778812001247

[7] A. Standard, "Standard 55-2017," *Thermal environmental conditions for human occupancy*, 2017.

[8] C. C. Federspiel and H. Asada, "User-adaptable comfort control for hvac systems," in *1992 American Control Conference*, June 1992, pp. 2312–2319.

[9] M. Feldmeier and J. A. Paradiso, "Personalized hvac control system," in *2010 Internet of Things (IOT)*, Nov 2010, pp. 1–8.

[10] V. L. Erickson and A. E. Cerpa, "Thermovote: participatory sensing for efficient building hvac conditioning," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 9–16.

[11] P. X. Gao and S. Keshav, "Spot: a smart personalized office thermal control system," in *Proceedings of the fourth international conference on Future energy systems*. ACM, 2013, pp. 237–246.

[12] F. Jazizadeh and S. Pradeep, "Can computers visually quantify human thermal comfort?: Short paper," in *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, ser. BuildSys '16. New York, NY, USA: ACM, 2016, pp. 95–98. [Online]. Available: http://doi.acm.org/10.1145/2993422.2993571

[13] F. Jazizadeh, A. Ghahramani, B. Becerik-Gerber, T. Kichkaylo, and M. Orosz, "User-led decentralized thermal comfort driven hvac operations for improved efficiency in office buildings," *Energy and Buildings*, vol. 70, pp. 398–410, 2014.

[14] P. Mansourifard, F. Jazizadeh, B. Krishnamachari, and B. Becerik-Gerber, "Online learning for personalized room-level thermal control: A multi-armed bandit framework," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 2013, pp. 1–8.

[15] F. Jazizadeh, F. M. Marin, and B. Becerik-Gerber, "A thermal preference scale for personalized comfort profile identification via participatory sensing," *Building and Environment*, vol. 68, pp. 140–149, 2013.

[16] F. Jazizadeh, A. Ghahramani, B. Becerik-Gerber, T. Kichkaylo, and M. Orosz, *Personalized Thermal Comfort-Driven Control in HVAC-Operated Office Buildings*. [Online]. Available: https://ascelibrary.org/doi/abs/10.1061/9780784413029.028

[17] F. Jazizadeh and B. Becerik-Gerber, "Toward adaptive comfort management in office buildings using participatory sensing for end user driven control," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2012, pp. 1–8.

[18] D. Daum, F. Haldi, and N. Morel, "A personalized measure of thermal comfort for building controls," *Building and Environment*, vol. 46, no. 1, pp. 3 – 11, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132310001915

[19] Y. Tachwali, H. Refai, and J. E. Fagan, "Minimizing hvac energy consumption using a wireless sensor network," in *IECON 2007 - 33rd Annual Conference of the IEEE Industrial Electronics Society*, Nov 2007, pp. 439–444.

[20] S. R. West, J. K. Ward, and J. Wall, "Trial results from a model predictive control and optimisation system for commercial building hvac," *Energy and Buildings*, vol. 72, pp. 271–279, 2014.

[21] T. A. Nguyen and M. Aiello, "Energy intelligent buildings based on user activity: A survey," *Energy and buildings*, vol. 56, pp. 244–257, 2013.

[22] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng, "Duty-cycling buildings aggressively: The next frontier in hvac control," in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*. IEEE, 2011, pp. 246–257.

[23] N. Nassif, S. Kajl, and R. Sabourin, "Evolutionary algorithms for multi-objective optimization in hvac system control strategy," in *Fuzzy Information, 2004. Processing NAFIPS'04. IEEE Annual Meeting of the*, vol. 1. IEEE, 2004, pp. 51–56.

[24] M. Fountain, E. Arens, R. de Dear, F. Bauman, and K. Miura, "Locally controlled air movement preferred in warm isothermal environments," 1994.

[25] A. ANSYS, "Academic research," *Release 16.0, Help System, ANSYS FLUENT Theory Guide*, 2015.

[26] D. Park and F. Battaglia, "Effect of heat loads and ambient conditions on thermal comfort for single-sided ventilation," in *Building Simulation*, vol. 8, no. 2. Springer, 2015, pp. 167–178.

[27] J. Wang, S. Wang, T. Zhang, and F. Battaglia, "Assessment of single-sided natural ventilation driven by buoyancy forces through variable window configurations," *Energy and Buildings*, vol. 139, pp. 762–779, 2017.

[28] D. Park and F. Battaglia, "Application of a wall-solar chimney for passive ventilation of dwellings," *Journal of Solar Energy Engineering*, vol. 137, no. 6, p. 061006, 2015.

[29] C. E. Rasmussen, *Gaussian Processes for Machine Learning (Adaptive Computation And Machine Learning)*. The MIT Press, nov 2005. [Online]. Available: https://www.xarg.org/ref/a/026218253X/

[30] P. O. Fanger *et al.*, "Thermal comfort. analysis and applications in environmental engineering." *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.

[31] J. Abanto, D. Barrero, M. Reggio, and B. Ozell, "Airflow modelling in a computer room," *Building and Environment*, vol. 39, no. 12, pp. 1393 – 1402, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132304001167

[32] M. Abedi, F. Jazizadeh, B. Huang, and F. Battaglia, "Smart hvac systems—adjustable airflow direction," in *Workshop of the European Group for Intelligent Computing in Engineering*. Springer, 2018, pp. 193–209.

[33] N. Nassif, S. Kajl, and R. Sabourin, "Optimization of hvac control system strategy using two-objective genetic algorithm," *HVAC&R Research*, vol. 11, no. 3, pp. 459–486, 2005. [Online]. Available: https://www.tandfonline.com/doi/abs/10.1080/10789669.2005.10391148

[34] S. Purdon, B. Kusy, R. Jurdak, and G. Challen, "Model-free hvac control using occupant feedback," in *Local Computer Networks Workshops (LCN Workshops), 2013 IEEE 38th Conference on*. IEEE, 2013, pp. 84–92.

[35] L. A. Hang-yat and D. Wang, "Carrying my environment with me: A participatory-sensing approach to enhance thermal comfort," in *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*. ACM, 2013, pp. 1–8.

[36] A. H.-y. Lam, Y. Yuan, and D. Wang, "An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings," in *Proceedings of the 5th international conference on Future energy systems*. ACM, 2014, pp. 133–143.

[37] F. Jazizadeh and W. Jung, "Personalized thermal comfort inference using rgb video images for distributed hvac control," *Applied Energy*, vol. 220, pp. 829 – 841, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261918301740

[38] W. Jung and F. Jazizadeh, "Towards integration of doppler radar sensors into personalized thermoregulation-based control of hvac," in *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments*, ser. BuildSys '17. New York, NY, USA: ACM, 2017, pp. 21:1–21:4. [Online]. Available: http://doi.acm.org/10.1145/3137133.3137166

[39] ——, "Non-intrusive detection of respiration for smart control of hvac system," in *Computing in Civil Engineering 2017*. Reston, VA, USA: American Society of Civil Engineers, 2017, pp. 310 – 17.

[40] A. Ghahramani, G. Castro, S. A. Karvigh, and B. Becerik-Gerber, "Towards unsupervised learning of thermal comfort using infrared thermography," *Applied Energy*, vol. 211, pp. 41 – 49, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S030626191731601X

[41] Z. Wu, N. Li, H. Cui, J. Peng, H. Chen, and P. Liu, "Using upper extremity skin temperatures to assess thermal comfort in office buildings in changsha, china," *International Journal of Environmental Research and Public Health*, vol. 14, no. 10, 2017. [Online]. Available: http://www.mdpi.com/1660-4601/14/10/1092

[42] N. Nassif, S. Kajl, and R. Sabourin, "Evolutionary algorithms for multi-objective optimization in hvac system control strategy," in *IEEE Annual Meeting of the Fuzzy Information, 2004. Processing NAFIPS '04.*, vol. 1, June 2004, pp. 51–56 Vol.1.

[43] K. Fong, V. Hanby, and T. Chow, "Hvac system optimization for energy management by evolutionary programming," *Energy and Buildings*, vol. 38, no. 3, pp. 220 – 231, 2006. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778805000939

[44] S. Wang and X. Jin, "Model-based optimal control of vav air-conditioning system using genetic algorithm," *Building and Environment*, vol. 35, no. 6, pp. 471 – 487, 2000. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132399000323

[45] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse, "The smart thermostat: Using occupancy sensors to save energy in homes," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 211–224. [Online]. Available: http://doi.acm.org/10.1145/1869983.1870005

[46] B. Dong and B. Andrews, "Sensor-based occupancy behavioral pattern recognition for energy and comfort management in intelligent buildings," in *Proceedings of building simulation*, 2009, pp. 1444–1451.

[47] B. Dong and K. P. Lam, "A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting," in *Building Simulation*, vol. 7, no. 1. Springer, 2014, pp. 89–106.

[48] X. Zhang, G. Schildbach, D. Sturzenegger, and M. Morari, "Scenario-based mpc for energy-efficient building climate control under weather and occupancy uncertainty," in *Control Conference (ECC), 2013 European*. IEEE, 2013, pp. 1029–1034.

[49] A. Parisio, D. Varagnolo, M. Molinari, G. Pattarello, L. Fabietti, and K. H. Johansson, "Implementation of a scenario-based mpc for hvac systems: an experimental case study," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 599 – 605, 2014, 19th IFAC World Congress. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474667016416800

[50] A. Beltran and A. E. Cerpa, "Optimal hvac building control with occupancy prediction," in *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, ser. BuildSys '14. New York, NY, USA: ACM, 2014, pp. 168–171. [Online]. Available: http://doi.acm.org/10.1145/2674061.2674072

[51] D. Gyalistras, A. Fischlin, M. Morari, C. Jones, F. Oldewurtel, A. Parisio, F. Ullmann, C. Sagerschnig, and A. Gruner, "Use of weather and occupancy forecasts for optimal building climate control," *Technical Report*, 2010.

[52] T. Sookoor and K. Whitehouse, "Roomzoner: occupancy-based room-level zoning of a centralized hvac system," in *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems*. ACM, 2013, pp. 209–218.

[53] Y. Chen, L. K. Norford, H. W. Samuelson, and A. Malkawi, "Optimal control of hvac and window systems for natural ventilation through reinforcement learning," *Energy and Buildings*, vol. 169, pp. 195 – 205, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778818302184

[54] S. Baghaee and I. Ulusoy, "User comfort and energy efficiency in hvac systems by q-learning," in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, May 2018, pp. 1–4.

[55] T. Wei, Y. Wang, and Q. Zhu, "Deep reinforcement learning for building hvac control," in *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*, June 2017, pp. 1–6.

[56] P. Fazenda, K. Veeramachaneni, P. Lima, and U.-M. O'Reilly, "Using reinforcement learning to optimize occupant comfort and energy usage in hvac systems," *Journal of Ambient Intelligence and Smart Environments*, vol. 6, no. 6, pp. 675–690, 2014.

[57] C. Huizenga, S. Abbaszadeh, L. Zagreus, and E. Arens, "Air quality and thermal comfort in office buildings: Results of a large indoor environmental quality survey," in *Healthy Buildings 2006*, 2006, pp. 393 – 397.

[58] M. Hamdi and G. Lachiver, "A fuzzy control system based on the human sensation of thermal comfort," in *Fuzzy Systems Proceedings, 1998. IEEE World Congress on Computational Intelligence., The 1998 IEEE International Conference on*, vol. 1. IEEE, 1998, pp. 487–492.

[59] L. Magnier and F. Haghighat, "Multiobjective optimization of building design using trnsys simulations, genetic algorithm, and artificial neural network," *Building and Environment*, vol. 45, no. 3, pp. 739 – 746, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132309002091

[60] N. Li, G. Calis, and B. Becerik-Gerber, "Measuring and monitoring occupancy with an rfid based system for demand-driven hvac operations," *Automation in Construction*, vol. 24, pp. 89 – 99, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0926580512000283

[61] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, "Sentinel: Occupancy based hvac actuation using existing wifi infrastructure within commercial buildings," in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '13. New York, NY, USA: ACM, 2013, pp. 17:1–17:14. [Online]. Available: http://doi.acm.org/10.1145/2517351.2517370

[62] S. Pan, A. Bonde, J. Jing, L. Zhang, P. Zhang, and H. Y. Noh, "Boes: building occupancy estimation system using sparse ambient vibration monitoring," in *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2014*, vol. 9061. International Society for Optics and Photonics, 2014, p. 90611O.

[63] E. Z. Conceição, J. M. Gomes, and A. E. Ruano, "Application of hvac systems with control based on pmv index in university buildings with complex topology," *IFAC-PapersOnLine*, vol. 51, no. 10, pp. 20 – 25, 2018, 3rd IFAC Conference on Embedded Systems, Computational Intelligence and Telematics in Control CESCIT 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S2405896318305470

[64] A. Afram, F. Janabi-Sharifi, A. S. Fung, and K. Raahemifar, "Artificial neural network (ann) based model predictive control (mpc) and optimization of hvac systems: A state of the art review and case study of a residential hvac system," *Energy and Buildings*, vol. 141, pp. 96 – 113, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778816310799

[65] L. Ciabattoni, G. Cimini, F. Ferracuti, M. Grisostomi, G. Ippoliti, and M. Pirro, "Indoor thermal comfort control through fuzzy logic pmv optimization," in *2015 International Joint Conference on Neural Networks (IJCNN)*, July 2015, pp. 1–6.

[66] N. H. Wong and S. S. Khoo, "Thermal comfort in classrooms in the tropics," *Energy and Buildings*, vol. 35, no. 4, pp. 337 – 351, 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778802001093

[67] R. Maiti, "Physiological and subjective thermal response from indians," *Building and Environment*, vol. 70, pp. 306 – 317, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132313002552

[68] K. Konis and M. Annavaram, "The occupant mobile gateway: A participatory sensing and machine-learning approach for occupant-aware energy management," *Building and Environment*, vol. 118, pp. 1 – 13, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132317301233

[69] X. Chen, Q. Wang, and J. Srebric, "Occupant feedback based model predictive control for thermal comfort and energy optimization: A chamber experimental evaluation," *Applied Energy*, vol. 164, pp. 341 – 351, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261915015159

[70] Z. J. Schlader, "The human thermoneutral and thermal comfort zones: thermal comfort in your own skin blood flow," *Temperature*, vol. 2, no. 1, pp. 47–48, 2015.

[71] W. Liu, Z. Lian, Q. Deng, and Y. Liu, "Evaluation of calculation methods of mean skin temperature for use in thermal comfort study," *Building and Environment*, vol. 46, no. 2, pp. 478–488, 2011.

[72] T. Sakoi, K. Tsuzuki, S. Kato, R. Ooka, D. Song, and S. Zhu, "Thermal comfort, skin temperature distribution, and sensible heat loss distribution in the sitting posture in various asymmetric radiant fields," *Building and Environment*, vol. 42, no. 12, pp. 3984 – 3999, 2007, indoor Air 2005 Conference. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132306003696

[73] C. F. Bulcao, S. M. Frank, S. N. Raja, K. M. Tran, and D. S. Goldstein, "Relative contribution of core and skin temperatures to thermal comfort in humans," *Journal of Thermal Biology*, vol. 25, no. 1-2, pp. 147–150, 2000.

[74] J.-H. Choi, V. Loftness, and D.-W. Lee, "Investigation of the possibility of the use of heart rate as a human factor for thermal sensation models," *Building and Environment*, vol. 50, pp. 165–175, 2012.

[75] S. A. Al-Sanea and M. Zedan, "Optimized monthly-fixed thermostat-setting scheme for maximum energy-savings and thermal comfort in air-conditioned spaces," *Applied Energy*, vol. 85, no. 5, pp. 326 – 346, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261907001201

[76] Y. Murakami, M. Terano, K. Mizutani, M. Harada, and S. Kuno, "Field experiments on energy consumption and thermal comfort in the office environment controlled by occupants' requirements from pc terminal," *Building and Environment*, vol. 42, no. 12, pp. 4022 – 4027, 2007, indoor Air 2005 Conference. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360132306003738

[77] D. Nikovski, J. Xu, and M. Nonaka, "A method for computing optimal set-point schedules for hvac systems," in *REHVA World Congress CLIMA'13*, 2013.

[78] C. Lin, C. C. Federspiel, and D. M. Auslander, "Multi-sensor single-actuator control of hvac systems," 2002.

[79] T. Sookoor, B. Holben, and K. Whitehouse, "Feasibility of retrofitting centralized hvac systems for room-level zoning," *Sustainable Computing: Informatics and Systems*, vol. 3, no. 3, pp. 161–171, 2013.

[80] V. Smith, T. Sookoor, and K. Whitehouse, "Modeling building thermal response to hvac zoning," *ACM SIGBED Review*, vol. 9, no. 3, pp. 39–45, 2012.

[81] X. Xu and S. Wang, "An adaptive demand-controlled ventilation strategy with zone temperature reset for multi-zone air-conditioning systems," *Indoor and Built Environment*, vol. 16, no. 5, pp. 426–437, 2007.

[82] Keen Home Inc. (2018) Keen home. [Online]. Available: https://www.keenhome.io

[83] Flair Manufacturing Corporation. (2018) Flair: Smart vents and wireless thermostats. [Online]. Available: https://www.flair.co

[84] L. J. Lo and A. Novoselac, "Localized air-conditioning with occupancy control in an open office," *Energy and Buildings*, vol. 42, no. 7, pp. 1120 – 1128, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778810000344

[85] W. Wang, J. Chen, G. Huang, and Y. Lu, "Energy efficient hvac control for an ips-enabled large space in commercial buildings through dynamic spatial occupancy distribution," *Applied Energy*, vol. 207, pp. 305 – 323, 2017, transformative Innovations for a Sustainable Future – Part II. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261917308139