

## **UC Merced**

### **UC Merced Electronic Theses and Dissertations**

#### **Title**

Optimization Framework For Improved Comfort & Efficiency

#### **Permalink**

<https://escholarship.org/uc/item/3p37v8xh>

#### **Author**

Yadav, Ashish

#### **Publication Date**

2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA  
Merced

**Optimization Framework For Improved Comfort &  
Efficiency**

A thesis submitted in partial satisfaction  
of the requirements for the degree  
Master of Science in Electrical Engineering and Computer Science

by

**Ashish Yadav**

2019

© Copyright by  
Ashish Yadav  
2019

The thesis of Ashish Yadav is approved.

---

Stefano Carpin

---

Wan DU

---

Alberto E. Cerpa, Committee Chair

University of California, Merced

2019



*Dedicated to my wife, mother, father and my big brother*

# TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b> . . . . .	<b>1</b>
<b>2</b>	<b>Related Work</b> . . . . .	<b>7</b>
2.1	Occupancy Detection & Prediction . . . . .	8
2.2	Occupant Comfort . . . . .	10
<b>3</b>	<b>System Requirements</b> . . . . .	<b>13</b>
3.1	Human Thermal Comfort . . . . .	15
3.1.1	Conditions of the occupant study . . . . .	15
3.1.2	Comfort App Design . . . . .	16
3.2	Occupancy Detection & Prediction . . . . .	18
3.3	System Identification . . . . .	22
<b>4</b>	<b>Model Predictive Control</b> . . . . .	<b>26</b>
4.1	Optimization Constraints . . . . .	26
4.2	Objective function . . . . .	27
<b>5</b>	<b>System Implementation</b> . . . . .	<b>30</b>
<b>6</b>	<b>Case Study</b> . . . . .	<b>35</b>
6.1	Environmental Conditions . . . . .	35
6.2	Performance Metrics . . . . .	36
6.3	Model Accuracy . . . . .	38
<b>7</b>	<b>Results</b> . . . . .	<b>41</b>

7.1	Thermal Comfort Analysis . . . . .	41
7.2	Cost and Energy Analysis . . . . .	45
7.2.1	Cost Analysis —& Energy Analysis . . . . .	47
<b>8</b>	<b>Discussion . . . . .</b>	<b>49</b>
<b>9</b>	<b>Conclusion . . . . .</b>	<b>51</b>
	<b>References . . . . .</b>	<b>52</b>

## LIST OF FIGURES

1.1	Energy consumption in a building: 2017 (US) . . . . .	2
1.2	HVAC System Architecture . . . . .	3
2.1	OFFICE System Overview . . . . .	8
3.1	Comfort voting page (left) and feedback (right) . . . . .	14
3.2	Comfort Voting page in web format . . . . .	16
3.3	Drift strategy setpoints to save energy . . . . .	17
3.4	Actuated flow (CFM) as occupancy changes . . . . .	18
3.5	GridEye Sensor before deployment . . . . .	19
3.6	8x8 thermal array sensing an occupant . . . . .	20
3.7	Occupancy sensors before deployment . . . . .	22
5.1	OFFICE System Architecture . . . . .	31
5.2	OFFICE Grafana Node Map . . . . .	33
5.3	OFFICE Zone Temperature Plots . . . . .	34
6.1	3-D Model of building . . . . .	36
6.2	Temperature prediction versus ground truth . . . . .	37
6.3	Model error characteristics for each zone . . . . .	38
6.4	Zone 1: Temperature prediction versus ground truth . . . . .	39
7.1	Example of thermal comfort for all zones in WebCTRL . . . . .	42
7.2	Thermal preferences vote distribution for the users in our building. Values [-3,-2,-1,0,1,2,3] corresponds to [Cold, Cool, Slightly Cool, Neutral, Slightly Warm, Warm, Hot] in our comfort app. . . . .	44

7.3	Zone temperature variation across the day for all zones for a selected day	44
7.4	Votes distribution in our comfort app as a function of the time of the day. . . . .	46
7.5	Operation cost of OFFICE vs WebCTRL . . . . .	47
7.6	Energy consumption of OFFICE vs WebCTRL . . . . .	48

## LIST OF TABLES

3.1	MPC variables and descriptions . . . . .	24
7.1	Room thermal satisfaction survey report . . . . .	41
7.2	Quality of Thermal Comfort . . . . .	42
7.3	Savings . . . . .	45

## ACKNOWLEDGMENTS

First, I would like to thank my advisor Alberto Cerpa for his support and guidance throughout my graduate study. I am also thankful for the delightful company and help of my colleagues, Alex Beltran, Daniel Winkler and Claudia Chitu who demonstrated to me what an excellent graduate student is supposed to be. I would also like to express my gratitude to my parents, Shiv Ram Yadav and Neelam Yadav, and my brother, Manish Yadav for their loving support and I am thankful to my wife, Suhani Nagpal for her trust and believing in me and making this work possible and my time at UC Merced all the more enjoyable.

Specifically, I would like to acknowledge the work included in this thesis which were co-authored by myself, my peers, and my Advisor. All chapters are based on our recent submission to Sensys '19. Section 3.2 of Chapter 3 is based on the publication, OBSERVER [ECC14], with primary author Varik L Erickson who used to work in the same lab. Section 3.1 is based on the publication Thermovote [EC12]. And once again, Professor Cerpa and Daniel Winkler who helped make all the above possible with guidance and editing.

I would like to thank and show my appreciation to all people involved in this project, for their volunteer time and valuable participation: from the personnel working in the building where this project took place to the assistant director from Facilities Management Team, James Brugger for his unconditioned support and expertise. Special thanks to Alex Beltran for his very important contributions and collaboration.

## VITA

1989	Born, Himachal Pradesh, India	2007-2011
2012-2017	Technical Associate, SpurTree Technologies–Bangalore,India.	2017-2019

## PUBLICATIONS

**Ashish Yadav**, Ahmed Sabbir Arif, Effects of Keyboard Background on Mobile Text Entry, In Proceedings of the Effects of Keyboard Background on Mobile Text Entry(MUM 2018) Nov 2018.



## ABSTRACT OF THE THESIS

# Optimization Framework For Improved Comfort & Efficiency

by

**Ashish Yadav**

Master of Science in Electrical Engineering and Computer Science

University of California, Merced, 2019

Professor Alberto E. Cerpa, Chair

The Internet of things (IoT) is the extension of Internet connectivity into physical devices and everyday objects. Embedded with electronics, Internet connectivity, and other forms of hardware (such as sensors), these devices can communicate and interact with others over the Internet, and they can be remotely monitored and controlled [Wik19].

The term internet of things has evolved by many fold due to convergence of multiple technologies like machine learning, embedded systems, wireless-sensors, real-time analytics, etc. A growing number of IoT devices are being developed for consumer use, including connected vehicles, home automation, wearable technology, connected health, and remote monitoring devices. We want to apply some of these technologies and techniques to make commercial spaces smarter.

Buildings are responsible for a significant portion of energy consumption in the US, accounting for more than 40% of US primary energy consumption. Heating, ventilation and air-conditioning (HVAC) accounts for nearly 50% of that use. Conditioning buildings is important since people spend 87% of their time in the place they live (residential) and the place they work (commercial). Despite this massive expense, many users are dissatisfied with the thermal conditions in buildings. Savings

made in HVAC systems therefore have a major impact in energy consumption and cost, together with reduction of greenhouse emissions for the nation. Equally critical is to provide thermal quality of service to their users, so people are comfortable in the place they reside and work.

In this thesis, we explore the trade off between commercial building HVAC energy consumption and the quality of thermal conditioning provided to users. We argue that optimal HVAC control cannot be achieved due to lack of critical information, namely *where* the users are inside the building, *what* do they want with respect to thermal comfort and *how* each zone responds to thermal changes. In this work we present OFFICE, a model predictive control (MPC) framework for smart building HVAC control. The framework has several components that help addressing the current HVAC control systems shortcomings, including

1. occupancy sensing in real-time
2. occupancy prediction models based on historical occupancy data
3. human-in-the-loop comfort feedback
4. data-driven thermodynamic building models, and
5. weather forecasting data

All these components provide the necessary input to our model predictive control optimization framework that minimizes monetary costs in energy use while maintaining quality comfort bounds for the building's users based on real-time user's feedback. We developed a large system that involves all the above components, replacing the Building Management System control algorithms, taking over full control of the HVAC system. We tested our framework OFFICE in a real LEED Gold certified university building with over 20 workers performing their daily tasks for 4 weeks, and we showed that we could obtain monetary costs savings of more than 10% while at the same

time reducing the users' dissatisfaction levels with thermal comfort from 25% to 0% dissatisfaction, significantly improving the quality of thermal service provided to the building's users.

# CHAPTER 1

## Introduction

Buildings are an essential piece of our daily lives and individuals invest 87% of their energy inside buildings [KNO01]. To keep up the thermal comfort in buildings, a lot of vitality is utilized to condition these spaces. In the US buildings represent 40% of energy utilization [Adm] and of that half of energy goes to heating, ventilation, and cooling (HVAC) [D12] Figure 1.1.

The goal of conditioning office spaces is largely missed since 75% of occupants report that they are dissatisfied with their thermal comfort [EC12]. In addition to uncomfortable occupants, a common issue is that spaces are conditioned whilst they may be unoccupied, or not ventilated appropriately based on the real quantity of occupants inside the room, losing enormous energy.

In the US, 88% of the large commercial building stock has centralized HVAC systems [D12]. The most prevalent HVAC system is the single duct terminal reheat, which is composed by an Air Handler Unit (AHU) that is a big fan (usually located on the building's roof) and heating and cooling coils that can modify the air's temperature, and Variable Air Volume (VAV) boxes that take this pre-conditioned air from the main duct, heat it if necessary (thus the terminal reheat), and control the air flow provided to its controlled zone. The air is mixed in each zone and excess air is returned by the return duct (which may have an additional fan). In many installations, the returned air goes to an economizer, which decides to recirculate a certain percentage of the air in order to save energy. Figure 1.2 shows the basic HVAC system architecture.

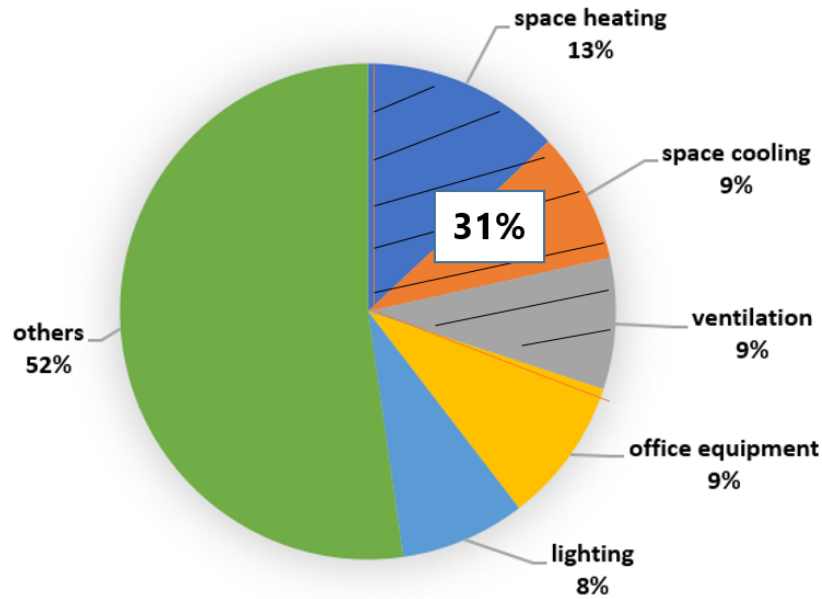


Figure 1.1: Energy consumption in a building: 2017 (US)

It is vital to point out that the principal intention of HVAC systems is to provide thermal comfort to its occupants, and doing it in an energy efficient way that minimizes power cost. So there is a clear trade-off in this situation, where in order to satisfy user thermal comfort requirements, the system must spend energy, which we ultimately want to minimize. If the only goal is to reduce energy consumption, the solution to this problem is simple, since we would simply turn off the HVAC system completely. We would also like to point out that by controlling both air flow and air temperature, the VAV achieves thermal comfort in each zone. Given the initial conditions (zone temperature) and the zone's thermodynamic model, the VAV's objective is to achieve a potentially new zone temperature in a finite amount of time. This goal can be achieved in many ways, and the set of possible solutions can be depicted by a line in the air flow and temperature two-dimensional space, where any point in the line provides a viable solution to the problem. For example, if a zone is too cold for user comfort and we need to increase the zone temperature, a possible solution could be providing a small value of air flow, but with an air temperature much hotter




Figure 1.2: HVAC System Architecture

than the zone goal temperature, such that after the hot supply air provided and the zone's air is mixed, the zone's temperature will increase and reach our goal temperature. Another solution could be to supply air that is only slightly hotter than our goal temperature, but increase the air flow significantly, achieving the same result. To make matters a little more complicated, zones are also controlled for ventilation requirements, which means that the variable air flow is also restricted by minimum flow requirements to remove the buildup of  $\text{CO}_2$  in the zone.

In this thesis, we argue that we need three pieces of information currently *missing* in all HVAC systems to design a system that minimizes energy costs subject to comfort constraints:

- (a) the distribution of occupants in all the zones/rooms (both real-time and future),
- (b) the thermal preferences of the occupants, and
- (c) the thermodynamic model of each zone.

In other words, we need to know *where* users are (and will be), *what* users want, and *how* each zone responds to changes. Currently, this information does not exist or is not available to HVAC engineers defining the operational control rules of the HVAC

systems. The state-of-the-art in building control is to make certain assumptions to fill this gap. First, occupancy is assumed to be homogeneous in all zones to maximum occupancy, so adequate ventilation levels can be set during the building “occupied hours”. Second, what constitutes a comfortable temperature range is defined by the facility manager in an *arbitrary* way during “occupied hours”, and those bounds are relaxed during the rest of the time (“unoccupied hours”). Finally, the HVAC control rules do not take advantage of any knowledge of the zone dynamics (e.g. loads, external walls insulation, solar gain, etc.), and they assume that the VAV boxes have been properly sized and configured at commissioning time when the building was constructed, and no changes to walls or internal space configuration have happened. Clearly, many of these assumptions are violated in practice, which leads to users being uncomfortable and building energy bills being higher than they should be.

It is not enough to know whether a room is occupied or not as previously explained in [ECC11], but rather it is critical to also know the *actual* number of occupants. This is because we can control the ventilation rate if we know the actual occupancy count based on ASHRAE 62.1 standard [ASH07b] and we can provide better thermal conditioning strategies with this information. In addition, since the time it takes to condition a room may be orders of magnitude larger than the time it takes users to move from one zone to another one, we need to *predict* the movement of building’s users in order to do pre-conditioning ahead of time. In our work, we leverage previous work in occupancy sensing [BEC13] and occupancy prediction [ECC14] to solve the answer to the question *where* users are (and will be).

In addition, we also leverage previous work related to human-in-the-loop feedback [EC12,WBE16] in order to get real-time estimation of users comfort requirements to address the question *what* users want. Note that by catering to *any* comfort needs of *any* user, we expect our solution to satisfy their quality of service requirements at the expense of potentially increasing energy use and cost.

Finally, we developed an online data-driven system identification process that

constantly collects data from building to find the optimum parameters of a gray-box model approach based on physical thermodynamic principles and model parameters from the data to answer the question of how each zone responds to changes.

OFFICE integrates all the above data together with weather prediction data into a MPC framework. Our system minimizes the cost of HVAC energy subject to the physical limitations of the HVAC system, the thermodynamic model learned from the system identification procedure, the user’s real-time and predictive occupancy, the user’s real-time comfort feedback and the current and predictive weather to achieve the best energy efficiency cost subject to the thermal comfort constraints of the user.

We would like to highlight the main contributions of our work as follows:

1. We develop a novel model predictive control (MPC) framework that optimally manages the trade off between energy cost and quality of comfort to the building users, by including input data from *where* users are (and will be), *what* users want, *how* zones react to changes, and current and forecast weather data.
2. We develop a novel data-driven system identification process together with a gray-box thermodynamic model, that it is simple enough to run in a control loop, yet powerful enough to permit sensible control decisions of the thermal zone.
3. We evaluate the system under realistic conditions in a LEED Gold certified building<sup>1</sup> [Cou17], which is being used by 20+ users in their daily tasks. We compare our results with the best control strategies designed by experienced building HVAC engineers, showing a significant improvement in both quality of thermal comfort to the users and overall energy cost.

To our knowledge, this is the first work that combines occupancy sensing and

---

<sup>1</sup>LEED: an acronym for Leadership in Energy and Environmental Design. The U.S. Green Building Council’s LEED green building program is the preeminent program for the design, construction, maintenance and operations of high-performance green buildings in the U.S.



prediction, human comfort feedback, data-driven thermodynamic models and weather forecasting into a holistic framework for HVAC building control.

The rest of my thesis is organized as follows: Chapter 2 shows the related work, Chapter 3 explains the main components that feed into our framework. Chapter 4 explains in detail the model predictive control framework, and Chapter 5 discusses system implementation details including the building communication interface for sensing and actuation. Then, Chapter 6 explains the conditions of the experiments, and Chapter 7 presents the results of the performance evaluation. Chapter 8 discusses the results and limitations of the system, and finally Chapter 9 concludes.

## CHAPTER 2

### Related Work

The use of MPC for HVAC control has been explored in the literature before [KMB13a, MKD12, KB11, MMB15, SGM16, CAN14, SFC18]. MPC is a very popular model because it provides opportunities for optimal energy management in HVAC control, being suitable in situations of conflicting constraints and objectives, such as physical constraints of the building, and indoor comfort ranges. However, most of this work has been done from a theoretical control theory point of view, with the experimentation mostly done in simulation.

In addition to traditional MPC models, work has been done in [MB12, PVR13, PMV13] to build a stochastic model predictive control. These efforts focus on uncertainty of occupancy and thermal load predictions and minimization the time the system falls outside the comfort ranges while optimizing the energy use. None of these systems do occupancy detection but consider occupancy as part of the load of the system, even though greater savings can be achieved with demand-response ventilation when occupancy is taken in consideration. In our work, we use deterministic MPC, i.e., whenever we have probability distribution functions to model our system inputs and constraints, we take the most likely value instead of using the full distribution of values. However, as we mentioned above, we consider a wider range of inputs into our MPC framework than this related work, including real-time and predictive occupancy, real-time user comfort feedback, data-drive thermodynamic zone models and weather forecasting to solve the optimization problem. While we could have used the more powerful stochastic MPC framework in our formulation, we decided to leave

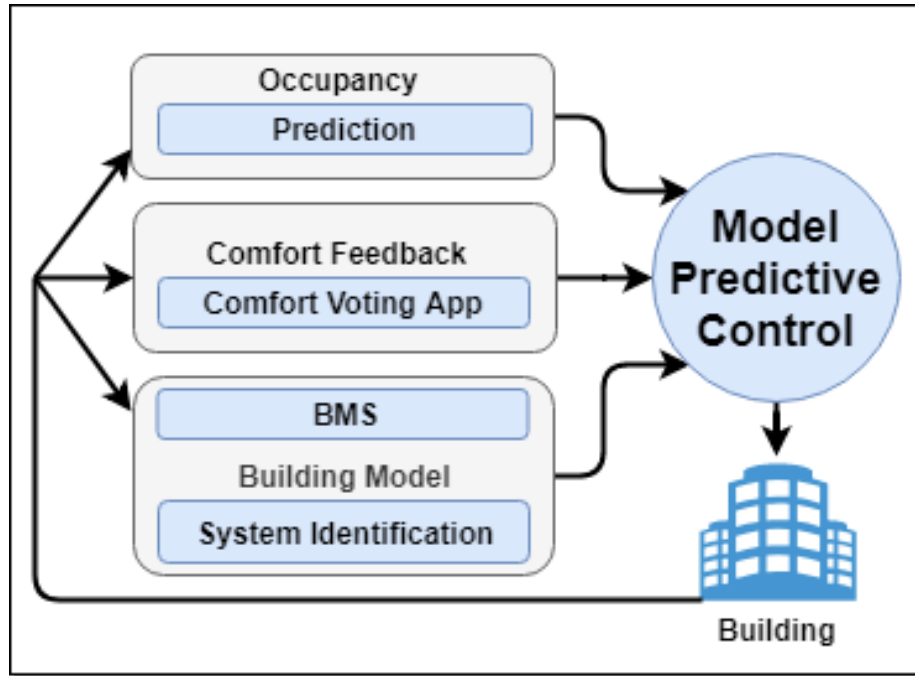


Figure 2.1: OFFICE System Overview

this for future work.

## 2.1 Occupancy Detection & Prediction

Due to its importance for the MPC, suitable approaches for occupancy estimation and forecast were studied extensively [ECC11, MCM12, EAC13, BEC13, ECC14, SAS17].

In [KJD09], the authors explore the use of cameras to estimate occupancy. They do not, however, address how cumulative error can affect estimates of occupancy. Even a single error will cause error to propagate forward, as mentioned above. The total ground truth for different times of the day was also limited to a total of 4 hours.

Cameras are also used as optical turnstiles by the authors of [EAC]. In order to measure occupancy for several areas, they mount multiple strategically placed cameras in hallways. Unlike [KJD09], however, they discuss cumulative error reduction strategies. To estimate the error, they impose maximum occupancy limits and use a

particle filter with a live data occupancy model. However, since their approach uses a model, their approach also requires a non-trivial amount of ground truth occupancy data (2 weeks), collected using webcams. The occupancy RMSE achieved was 1.83 persons, more than 5 times the ThermoSense’s [BEC13] error.

In [BLT10], occupant counts are achieved by counting peaks within the histograms, as well as PIR sensors in order to detect occupancy for certain areas and elderly people are tracked using Imote2 motes with Enalab cameras and utilizing a motion histogram for a period of 1 week. Since the camera has to continually poll the room to generate the histograms of motion, the power consumption will be high during occupancy periods. This system also has the privacy issue ; cameras need to be placed directly in the room.

The authors of [SBK11] use active radio frequency identification (RFID) tags to determine occupancy instead of relying on PIR sensors alone. The limitation of this strategy is that each occupant must have an RFID tag and that tag must be co-located with the occupant at all times.

The papers [LSS10, ABD11] describe methods that uses door sensors with PIR sensors to obtain a binary measurement of occupancy. They minimize instances in which overly still occupants become invisible to the PIR sensor by adding door sensors. Although this technique improves binary occupancy measurement, these systems do not provide an accurate estimate of occupancy.

The papers [LHD09, MMC11] estimates occupancy by measuring a variety of parameters. For 5 and 1 week deployments, they collected ground-truth data using video camera and a voluntary electronic tally counter for the user to measure room occupancy. In these deployments, they utilize multiple sensors to estimate occupancy; CO<sub>2</sub>, CO, lighting, temperature, humidity, motion, and acoustics. They define multiple feature vectors for each parameter, which are then used to estimate occupancy with multiple models. While this multi-sensor approach works well for occupancy estimates alone, when combined with a ventilation strategy, this approach will not work

well. They assume that room occupancy estimates will not be affected by ventilation. However, as ventilation will affect CO<sub>2</sub> and humidity levels and therefore occupancy estimates, ventilation rates based on occupancy estimates from this system are likely to result in wild fluctuations in ventilation actuation and underventilation periods. In this case, even with the known calibration and response time problems [FFS06], ventilation is better controlled directly by CO<sub>2</sub> sensors. Essentially, if CO<sub>2</sub> or humidity is used as a sensory input, either you can control ventilation or estimate occupancy, not both at once.

The work in [GIB12] was the first attempt to quantify the inclusion of occupancy data with an MPC control strategy, showing its great potential via simulations. This was followed by [BC14], which used a very specific occupancy sensor infrastructure to provide real-time occupancy in an MPC formulation to provide additional energy savings. In our approach, we take into consideration not only real-time occupancy, but predicted occupancy as well. Perhaps more importantly, our work includes real-time user’s comfort feedback as well as a data-driven thermodynamic model of each zone to guide the MPC optimization process. So in our work, we use the model developed in [ECC11, ECC14] in order to predict future occupancy based on real-time data.

## 2.2 Occupant Comfort

Thermal comfort is critical as the main goal of an HVAC system is to provide quality of thermal comfort to the building users. However, it is difficult to establish objective models of comfort that represent human preferences. This issue has been debated for a long time because of its somehow subjective definition. The first work addressing this issue was Fanger’s seminal work on Predictive Mean Vote (PMV) [Fan70]. However, this model is subject to significant differences observed across gender, age, clothing preferences, physical fitness, as recent studies and experiments in built environments have proved it recently [RVL15, WPL19]. In order to cope with this model limitations,

many systems have been designed to include the human-in-the-loop and ask users to provide their real-time thermal preferences.

Previous work have also used building occupants as participatory sensors [RES10]. These systems collect thermal comfort information from building occupants by bringing humans into the loop [EC, BTG13, JB12, Bur14, Rob14], and then use PMV to determine temperature setpoints to improve comfort. These works conclude that collecting sensation data like "Feeling Too Cold" improves usability because users can not determine their ideal temperature. All studies resulted in savings in energy and enhanced satisfaction. In addition to PMV, a multi-arm bandit framework [MJK13] and an optimization model [HW13] were used to identify enhanced temperature setpoints based on voting patterns.

Comfy [Rob14] is a system which also gives users the ability to vote for their comfort. They take votes in the form of Hot, Comfy, Cold. They were able to reduce HVAC cost by 22.95% by doing additional strategies such as expanded the heating and cooling setpoints to have a larger deadband. This is similar to our methods, but the algorithm used for saving energy is not stated.

Including human-in-loop is done mostly by a voting platform, as it has been the case of many systems in the literature [EC12, PKJ13, JGB13, WBE16]. Participatory sensing with a human-in-the-loop system has many advantages, including more precise information of specific users' thermal preferences, more democratic thermal consensus when multiple user sharing a zone have different thermal preferences, and overall energy efficiencies by conditioning the spaces based on what the users actually want instead of trying to guess the comfort range by a subjective decision taken by a facility manager. In our paper we include the feedback mechanisms to motivate a user to vote and to keep them engaged.

Research carried out in [YAL14] investigates the use of various feedback mechanisms to improve an occupant's office space's energy efficiency by reminding / enabling the occupant to disable unnecessary equipment (lights, computers, etc.) when the

user is not in the office. In our work, we wish to reduce the energy consumption of the HVAC system while it conditions the user's space. For this reason, we must leverage energy savings against occupant comfort. We included a physical feedback mechanism which a voter can feel as soon as he/she votes from the vents in the room.

In our work, we leverage these techniques and use the ThermoVote system [EC12] as our comfort app for OFFICE. Furthermore, we use the drift control strategy previously developed in [WBE16].

In summary, we believe our work is the first attempt in the literature to holistically combine real-time and predictive occupancy, real-time thermal user feedback, data-driven thermodynamic models and weather forecast into a model predictive control framework in order to optimally control the tradeoff between energy use and comfort.

## CHAPTER 3

### System Requirements

Figure 2.1 depicts a high-level overview of the our system. When the system is installed, occupants within the space are given access to a *thermal comfort voting application*, which allows each user to vote for their comfort in their working space as frequently as they like. In addition, a wireless network of low-power *occupancy sensors* are installed overhead to track the distribution of occupants across the space during operation. As the HVAC system runs its control strategy and daily operation occurs in the space, we collect votes from the occupants and track their movement using the overhead sensors. In addition, as the HVAC system runs, we periodically (every 2 minutes) sample the current values of temperature air flow in each zone and the main loop via the BACNet communication protocol, as later explained in Section 5. With this data collected, we have a clear understanding of the recent and current thermal conditions of the space, occupancy distribution, and comfort.

Every 5 minutes, our systems processing pipeline is started, which first retrieves all recent votes, occupancy distribution, and building data samples from the databases where they are stored. With the recent comfort votes from the users across the space, the system chooses the appropriate heating and cooling set points for each zone, such that the occupants will remain comfortable. Using the current distribution of occupants, the system identifies unoccupied or under-occupied zones, candidates for reduced conditioning for energy savings. Using a blended Markov chain [ECC11] trained on historical occupancy data from the running building, the current occupancy is sampled from the model to statistically predict the occupancy in the near



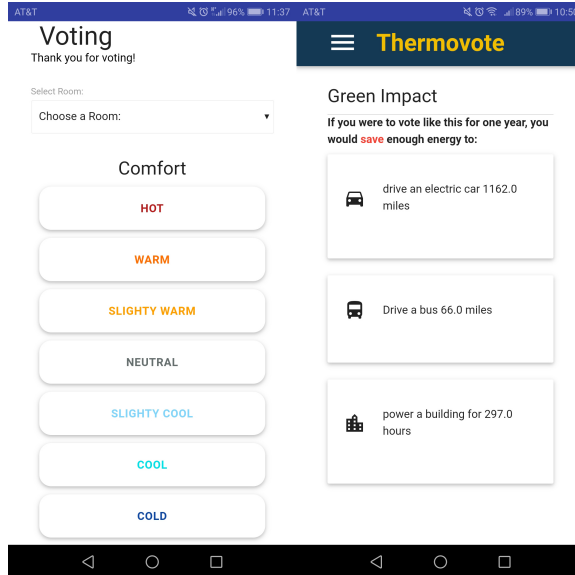


Figure 3.1: Comfort voting page (left) and feedback (right)

future. Finally, local weather prediction is queried from the public Wunderground API [wun19] to determine expected weather trends.

Finally, we maintain a gray-box thermodynamic model of the building using data-driven *system identification*, which will later be used to predict how the thermal state of the building will respond to a given control sequence. These collected data samples and models are all then used to define, initialize, and solve a model predictive control optimization problem that finds the building actuations minimizing system operation costs subject to occupant comfort constraints based on their voting feedback. These optimized actuations are then communicated with the building and applied for each zone and the main AHU loop. After sufficient time has passed, the processing pipeline will be run again, accounting for any new conditions that have been sensed in the space for future control.

## 3.1 Human Thermal Comfort

### 3.1.1 Conditions of the occupant study

To better understand the initial environment and the human factor particularities we run a pre-study survey. This study is approved by the IRB (Institutional Review Board), and the participants were selected to be those who work in the space chosen as the location for our predictive thermal control. The subjects in our study are from a large range of ages from 21 to 60 years old, with a 37% female population from a total of 16 people. From this study we found that 75% of them work between 31 and 40 hours per week from 6AM to 7PM, being exposed to different thermal conditions across the workday. 31.25% have a private office and the rest of them shares a space with few other colleagues. When asked how their thermal preferences compare with their colleagues, 31.25% of the volunteers answered that they prefer a warmer room and 31.25% specified that they prefer a colder room. 56.25% of the people that took this pre-survey answered that they have contacted the facilities team regarding the temperature conditions up to *8 times per person* in the last 12 months. 12.50% of them found their thermostat as ineffective and 75% of them do not have a thermostat to control the room temperature. 37.50% of them reported that room temperature enhances their ability to work and 25% believe that it interferes with their ability to work. 37.50% said that they avoided their office because of thermal discomfort and 80% of the total population did clothing adjustment (more or less clothes) when it was necessary. This supports the hypothesis that the thermal comfort has a great impact on the self reported productivity in office spaces, in addition to personal satisfaction. We asked questions related to a new control such as how quickly they would expect their space to be conditioned after adjusting the thermostat and 56.25% answered that they expect to be done in more than 10 minutes. Asking the study's participants about their satisfaction previous to the new building control, we found that 25% of them felt dissatisfied. With this image in mind, we designed a system to

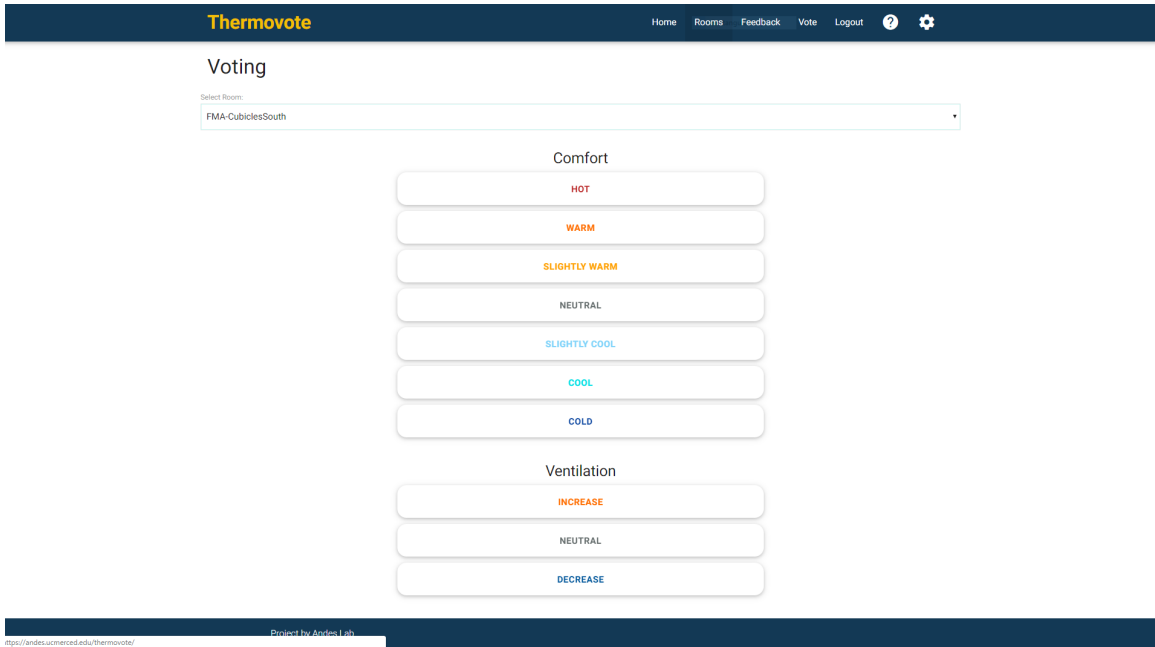


Figure 3.2: Comfort Voting page in web format

better suit the occupants’ preferences, which thrives in dynamic thermal conditions and is built on a modern infrastructure including open source frameworks.

### 3.1.2 Comfort App Design

To facilitate user interaction with our system, we designed an application in a web format(website) Figure 3.2 distributed to the occupants for iOS and Android platforms, with credentials. It is based on HTML5 and with this feature, content changes in one place are automatically propagated on the mobile phone versions, without passing through the AppStore or Google Play for updates. The app “look and feel” is shown in Figure 3.1. In the left side is the page with the vote buttons whilst in the right side is the page with the green feedback, designed to help the user understand their energy footprint in relatable terms, such as how many hours an electric car could travel or how many hours a building could be powered. Once a vote is issued by a user it is stored in a MySQL database, from where the system incorporates the average vote per zone to find the temperature bounds for that specific zone which will be input

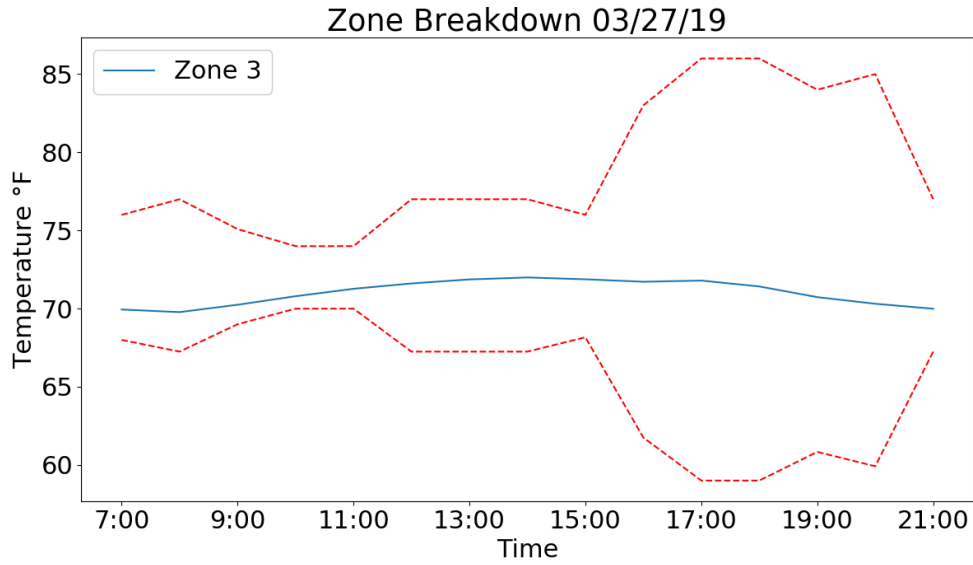


Figure 3.3: Drift strategy setpoints to save energy

into the MPC as constraints.

In our experiments, we allow the bounds to start drifting apart by  $1/2^\circ\text{F}$  every 30 minutes, but remaining between 68 and  $76^\circ\text{F}$  for occupied space and between 55 and  $90^\circ\text{F}$  for unoccupied zone, according to the policies chosen by the building managers. This drift begins when 1 hour has passed since the last issued vote in a zone and makes possible energy savings by allowing the temperature to “float” between these expanding bounds. In normal usage, heating/cooling setpoints are chosen based on occupant voting patterns. However, if these bounds are held unnecessarily tightly, extra energy will be consumed. An example of using drift strategy is shown in Figure 3.3. This figure illustrates that the bounds have started to relax at 12:00 when people go for lunch and in the vote’s absence, at 13:00 they reached the occupied limits (provided by the facility management team) and after 17:00, since the building was empty, the bounds exceeded the previous limits to go closer to the ones for the unoccupied schedule. The zone temperature is not correlated with the bounds values but it is important to remain within the limits. The way the bounds are changing with the drift is retrieved from [WBE16], using the constants mentioned there for the




Figure 3.4: Actuated flow (CFM) as occupancy changes

Fanger formulation (PMV) [Fan70]. In practice, we receive votes from the users that describe how they feel on a scale from cold to hot as in Figure 3.1. By taking the average of all recent votes in the zone as an actual mean vote (AMV) and setting equal to the PMV formulation, we find new bounds values for use in control.

### 3.2 Occupancy Detection & Prediction

In order to truly condition a building efficiently, it is necessary to be able to reliably detect the distribution of occupants throughout the space. A system can save significant energy based on occupancy, as both ventilation and conditioning codes



Figure 3.5: GridEye Sensor before deployment

are explicitly defined based on occupancy levels as defined by ASHRAE standards [ASH07b] [ole73]. In an unoccupied zone, ventilation can be set to a minimum level based on the square footage of the zone, and as occupancy increases, ventilation must increase proportionally to prevent buildup of  $\text{CO}_2$  and other undesired contaminants produced by the occupants as demonstrated in Figure 3.4. We can see in the figure that as the mean occupancy increases, the ventilation tends to closely follow, and that when occupancy goes to zero, ventilation holds steady at a minimum level based on the square footage to be code compliant. Thermal conditioning of the air exhausted into the space, however, changes based on the 0/1 boundary of occupancy. If there are no occupants in the space, the air blown into the zone does not need to be conditioned, saving a significant amount of energy. However, if at least one occupant is present, the temperature of the zone must be maintained to keep the occupants comfortable within the space.

For the our system to consider occupancy in its control decisions, we wish to have both occupancy measurement *and* prediction for the building. For occupancy mea-



Figure 3.6: 8x8 thermal array sensing an occupant

surement, we employ a wireless sensor network of occupancy sensors with Panasonic PIR and Grideye sensors as proposed by [BEC13] and shown in Figure 3.7 and Figure 3.5. The installed devices periodically poll the PIR sensor for recent movement, and if necessary the Grideye captures a thermal image Figure 3.6. Using an exponential weighted moving average to maintain a background of the space, background subtraction is used to find active regions in a new frame based on the standard deviation of temperature values for each pixel. Then, running an 8-connected components algorithm, we identify the distinct active “blobs” in the frame, which identify the unique heat sources in the frame. As demonstrated in [BEC13], we use a linear regressor trained on historical data that maps the size of the largest active blob, the number of blobs, and the number of total active pixels in the image to the predicted number of people in the frame, which we have found to work well in practice. The number of people output by the regressor is then stored in a database to be used later in the processing pipeline.

To allow occupancy *prediction*, we use a blended markov chain as introduced in [ECC11]. This model is trained with historical data traces of occupancy in this building, and maintains transition matrices for occupancy distribution specific to the time-of-day. This allows our system to query the model with the current distribution of occupants throughout the space and get the predicted future sequence of occupancy, with states in the model occurring closest to the current time-of-day being most heavily considered. Since we only consider states observed in the training data as opposed to observable states, there might be states that are not available in a transition matrix for that specific time of day. Suppose we are predicting occupancy for hour  $H_x$ . After 3600 steps, we are in some state  $S$ , the hours changes from  $H$  to  $H_x$ , and the model will switch to  $H_x$  hourly transition matrix. It is possible the transition matrix for hour  $H_x$  for occupancy state  $S$  has no probability. This happens when state  $S$  never occurs in the  $H_x$  training data. Even though in reality the state  $S$  can occur in hour  $H_x$ , if  $S$  does not occur in hour  $H_x$  of the training data, then we cannot calculate the transition probabilities for  $S$ . This can not be solved by introducing a small probability value of transition to another state as the next selected state may not be represented in the transition matrix as well. The markov chain becomes a random walk until the observed data set captures a particular state. Also, sink state can occur if the next transition matrix does not contain a probability for the current state. To tackle these issues, rather than considering the closest distance transitions we linearly combine the different transition matrices to obtain blended transition matrix to include all observed states. This increases the number of preferred states available for transition and decreases the chance to select states outside the slot boundaries completely.

Specifically, this type of model will learn time-specific behaviors and sequences, such as occupants morning influx when business hours start, or noontime lunch movement. With the ability to predict future occupancy, our system will be able to consider future movement of people and pre-condition the space accordingly.





Figure 3.7: Occupancy sensors before deployment

### 3.3 System Identification

In optimization, we must consider how the building will react to system inputs such as zone flow and temperature adjustment. As shown in [KMB13b], a resistor-capacitor model can be used to model the thermodynamics of a building where the capacitance represents the heat storage capacity of each zone, analogous to room size, and resistance represents how insulated the zones are from each other (i.e. thickness of walls/doors). In this way, we model zone temperature dynamics as follows with model variables enumerated in Table 3.1:

$$M_n \frac{d}{dt} T_{z_n} = (T_{s_n} - T_{z_n}) c_p m_{z_n} + Q_n + (T_o - T_{z_n}) / R_{n,o} \quad (3.1)$$

Here, the three terms on the right side represent heat flow from the HVAC system through the vents into the zone, the load due to equipment or machinery, and heat flow from the outside of the building into the zone, respectively. In addition, the  $M_n$  term on the left hand side denotes the capacitance of the zone, impacting the rate at which the zone's thermal status can be changed. While there may be heat exchange *between* zones as well, temperature conditions between two indoor zones will

very similar in general, making the heat transfer negligible. In addition, by including these terms, the model becomes highly non-linear, making optimization significantly more difficult.

Before this thermal model can be used, we must estimate values for each zone’s capacity ( $M$ ), thermal load ( $Q$ ), and thermal resistance to the outside of the building ( $R$ ). In practice, the load within each zone and resistance to the outside are difficult to estimate accurately by simply observing the space; factors that contribute to load may be difficult to observe, such as electronic equipment whose heat varies over time depending on usage, appliances in the space that are not constantly running, etc. Similarly, the resistance to the outside of the building is dependent on the insulative abilities of the walls, windows, and doors of the external walls and the area each covers the wall. In contrast, the capacitance of each zone in the building will be directly proportional to the amount of air that can be stored in the zone, which can be determined through most standard floor plans of the indoor space. We note that while this estimate will not consider the space occupied by furniture and people, we consider it a fairly close approximation. In this way, we compute the  $M$  values directly, reducing the number of unknown parameters that must be fit from data in the next step.

The remaining parameters,  $R$  and  $Q$ , are learned periodically using thermal data traces collected within the building. In this way, our model will maintain its relative accuracy over time, as these conditions may be changing based on equipment and personnel usage of the space. To estimate the resistance and load parameters of our system, we use our thermodynamic model of Equation 3.1 and data traces of the system in operation in a least squares analysis to find the values of  $Q$  and  $R$  that minimize the error of the model given those parameters. This operation is run daily in our system to ensure the model will adapt to changing conditions in the space, such as movement of equipment between zones (printers, computers) or changing working hours of the occupants.

Type	Symbol	Description
Constants	$M_n$	Capacity of a zone
	$c_p$	Heat capacity of air
	$R_{n,i}$	Resistance between zone $i$ and zone $n$
	$R_{n,o}$	Resistance between zone $n$ and outside
	$\eta_c, \eta_h$	Cooling and heating efficiency
	$r_g$	Monetary cost of gas per kwh
	$r_e$	Monetary cost of electricity per kwh
	$\tau, N$	Number of time steps and zones
	$T_{c,min}(t)$	Min cooling temperature
	$T_{s,max}(t)$	Max heating temperature
	$R_o$	Required ventilation per occupant
	$R_{a_n}$	Required ventilation for an area
	$A_n$	Area of a zone
Variable	$T_{z_n}(t)$	Zone Temperature
	$T_{s_n}(t)$	Supply air temperature to a zone
	$T_c(t)$	Cooled air temp from AHU
	$T_m(t)$	Mixed Air Temperature
	$D(t)$	Damper position for outside air
	$T_r(t)$	Return temperature
	$m_{z_n}(t)$	Mass flow into a zone
	$m_s(t)$	Total mass flow of all zones
	$\rho_{u,n}(t), \rho_{l,n}(t)$	Upper/lower comfort penalty
	$\phi_p$	Penalty coefficient
	$t, n$	Time, Zone
Given or Pre-calculated	$T_o(t)$	Outside air temperature
	$Q$	Thermal load
	$O_n(t)$	Occupancy
	$V_{min,n}(t)$	Min ventilation required for a zone
	$T_{z_n}^+(t), T_{z_n}^-(t)$	Upper and lower comfort bounds

Table 3.1: MPC variables and descriptions

As with any process where parameters must be fit to data, it is important that the datapoints used are well-distributed, representing the full data space [Ram69]. In other words, if all datapoints have the same or similar values, a bad parameter fit is likely. For example, if we use temperature values from a conditioned space, the value of load and resistance won't be reasonable. We use least squares regression to choose the best parameters for our chosen model using the data traces from the building. In our application, during building operation the indoor temperature conditions are relatively constant. For this reason, we require homogeneous distribution of input points on the entire temperature range. We have found that it is best to use data from the mornings when the HVAC system first conditions the space, and nights/weekends when the HVAC system is turned off and the interior temperature drifts due to load and external conditions only.

# CHAPTER 4

## Model Predictive Control

### 4.1 Optimization Constraints

To constrain the system, we first must include the thermodynamic model as introduced in Equation 3.1:

*Constraint 1:* Thermal Model of Equation 3.1

The remaining optimization constraints are based on the physical properties of the system, building regulations that the HVAC system must follow, and comfort constraints. In addition, in cases where zone temperatures are outside of comfort ranges as initial conditions for optimization, an upper and lower penalty is added as well to prevent the optimization from becoming infeasible, but guide the zone temperature into a comfortable range as quickly as possible.

*Constraint 2:*  $T_c(t) \leq T_m(t)$  – The cooling coil can only cool the air received from the economizer.

*Constraint 3:*  $T_c(t) \geq T_{c,max}$  – Cooling capacity of cooling coil.

*Constraint 4:*  $T_s(t) \geq T_c(t)$  – Heating coils can only increase the temperature of the air supplied by the AHU.

*Constraint 5:*  $T_s(t) \leq T_{s,max}$  – Max capacity of heating coil.

*Constraint 6:*  $m_{z_n}(t) \geq V_{min,n}(t)$  – The zone's minimum ventilation required for occupants.

*Constraint 7:*  $m_{z_n}(t) \leq m_{max_n}$  – The VAV’s maximum ventilation capacity.

*Constraint 8:*  $0 \leq D(t) \leq 1$  – Physical damper constraint.

*Constraint 9:*  $D(t) \geq D_{min}$  – Minimum outside air requirement.

*Constraint 10:*  $\sum_{n=1}^N m_{z_n}(t) \leq AHU_{max}$  – The fan’s maximum ventilation capacity.

*Constraint 11:*  $\rho_{u_n}(t) + T_{z_n}^+ \geq T_{z_n}(t) \geq T_{z_n}^- - \rho_{u_n}(t)$  – Comfort bound

*Constraint 12:*  $\rho_{u_n}(t) \geq 0, \rho_{l_n}(t) \geq 0$  – Penalty functions can only increase cost

By honoring these system constraints, the model predictive control will respect the requirements and limitations of the physical HVAC system when considering potential actuations.

## 4.2 Objective function

To choose the best actuation sequences possible, we create the objective function such that the cost of system operation is minimized. In our HVAC system, the three primary energy consumers are AHU ventilation and cooling (both electricity) and heating (gas). While there are further consumers such as water pumps for heating and cooling, these are more minor consumers, and are not considered in our model. Similarly to the thermodynamic model previously presented in Section 3.3, we can use the temperature values before and after heating and cooling to compute the amount of energy that has been transferred. For instance, we can compute the energy used for cooling at times  $t \in 1 \dots \tau$  as

$$P_c = \frac{c_p}{\eta_c} \sum_{t=1}^{\tau} m_s(t)(T_m(t) - T_c(t)) \quad (4.1)$$

where  $T_m(t)$  is the warmer mixed air,  $T_c(t)$  is the cooler output,  $m_s(t)$  is the mass flow rate,  $c_p$  is the heat capacity of air, and  $\eta_c$  is the efficiency of the air conditioner. With temperature units in Kelvin and mass flow in kilograms per second and other values as

defined in Table 3.1, this equation will provide the kWh consumed in heating, which can then be multiplied into the cost of electricity  $r_e$  to convert to monetary value. The mixed air temperature  $T_m$  is a mixture of the return temperature  $T_r$  and the outside temperature ( $T_o$ ), the ratio of which is controlled by the economizer damper position  $D(t)$ . At  $D(t) = 1$ , 100% of the mixed air is recirculated, and at  $D(t) = 0$ , 100% of the air comes from the outside.

$$T_m(t) = D(t)T_r(t) + (1 - D(t))T_o(t) \quad (4.2)$$

While fresh air is required to prevent buildup of gases in the space, more re-circulation allows the system to save energy, as the re-circulated air has already been conditioned to roughly the desired temperature. The return temperature, then, is a mixture of outgoing air temperature from each zone, weighted by that zone's mass air flow:

$$T_r(t) = \frac{\sum_{n=1}^N m_{z_n}(t)T_{z_n}(t)}{m_s(t)} \quad (4.3)$$

The technique to measure energy consumed for heating is similar to that of cooling. We consider the temperature increase across the heating coil in each zone and the mass flow of air to get

$$P_h = \frac{c_p}{\eta_h} \sum_{i=1}^N \sum_{t=1}^{\tau} m_{z_i}(t)(T_{s_i}(t) - T_c(t)) \quad (4.4)$$

where  $m_{z_i}$  is the mass flow for zone  $i$ ,  $T_{s_i}$  is zone  $i$ 's supply temperature into the room, and  $T_c(t)$  is the cool air temperature in the main loop. Finally, we consider the cost of the fan, which was found to be a squared increase with respect to mass flow in [KMB13b] with an efficiency factor of  $\eta_f$ :

$$P_f = \sum_{t=1}^{\tau} \eta_f m_s(t)^2 \quad (4.5)$$

Considering the differing costs for gas and electricity, the minimization problem can be written as follows where  $r_e$  and  $r_g$  are the dollar costs per kWh for electricity and gas, respectively. In our locale, these pricing figures are .018 \$/kWh for gas and

.108 \$/kWh for electricity. In our deployment, this means the optimized actuations will prefer to provide more gas heating due to the cheaper costs, but under different pricing conditions OFFICE will find the appropriate control that will minimize costs for any pricing.

$$\min_{\substack{m_n(t), T_{zn}(t), T_{sn}(t), T_c(t), D(t) \\ t \in 1 \dots \tau, n \in 1 \dots N}} r_e(P_c + P_f) + r_g(P_h) + C_s \quad (4.6)$$

subject to: Constraints 1-12

As the objective function and constraints are nonlinear with respect to the optimization variables, the optimization problem as written is non-convex which may cause the optimized solutions to be local and not global optima. Future work may consider the process of model linearization, but analysis must be done to ensure model accuracy is not substantially reduced. In practice, our system models the thermodynamics of the 8 zones of our test building ( $N = 8$ ) 1 hour into the future, at 10 minute intervals ( $\tau = 6$ ). Despite model non-convexity, a solution is found in seconds using the Julia [BEK14] optimization framework’s IPOPT solver for our test building.

When optimization is finished, it returns the optimal setpoints of discharge temperature and mass flow for each VAV in the system, damper position in the economizer, and temperature setpoint for cooling in the AHU for the next hour into the future at 10 minute intervals. In addition, for each of these discrete timesteps, the optimization provides the state of the system subject to the optimal control, as predicted by the thermodynamic model. In practice, although optimization considers the evolution of the system for a time window of 1 hour into the future, when optimization is finished only the first set of actuations (i.e. actuations at  $t = 0$ ) are sent to the building for actuation. Then, 10 minutes in the future, all data is queried and optimization is completed again to correct for any error in prediction that may have occurred.



## CHAPTER 5

### System Implementation

Our system has several modules that require data gathering capabilities, as discussed in Sections 5 and 4. In this section we explain the building interaction with various modules that form this system, and practical considerations in its design. Figure 5.1 provides a detailed view of the system architecture, along with the backbone storage systems that maintain the large amount of data required in processing. Due to the different types of data we maintain for our processing pipeline, we maintain two distinct databases; for data sources that are periodic such as building sensor data, a time-series data called Influx [inf13] is used, and for data that is not time-series is stored in a relational MySQL [MyS95] database.

Building control is not easy. Incorrect actuation commands can result in discomfort to the occupants in the space, reduction of hardware lifetime through an unexpectedly high number of physical cycles, or even direct damage of HVAC equipment. This is especially risky due to the reliance of the system on high data quality from the various data sources in the space. In our experience, it is common for embedded sensors within the building equipment to become miscalibrated or simply fail, and similarly we have found actuators that will not obey a request to change states. For these reasons, the highly modularized structure of our system has been instrumental to maintaining a high quality of control. To ensure system operation, a number of fail safes have been implemented; features within the Influx time-series database have allowed us to create custom alarms to sense when sensors have not provided data for a period of time, allowing us to detect and correct issues in our deployments, and

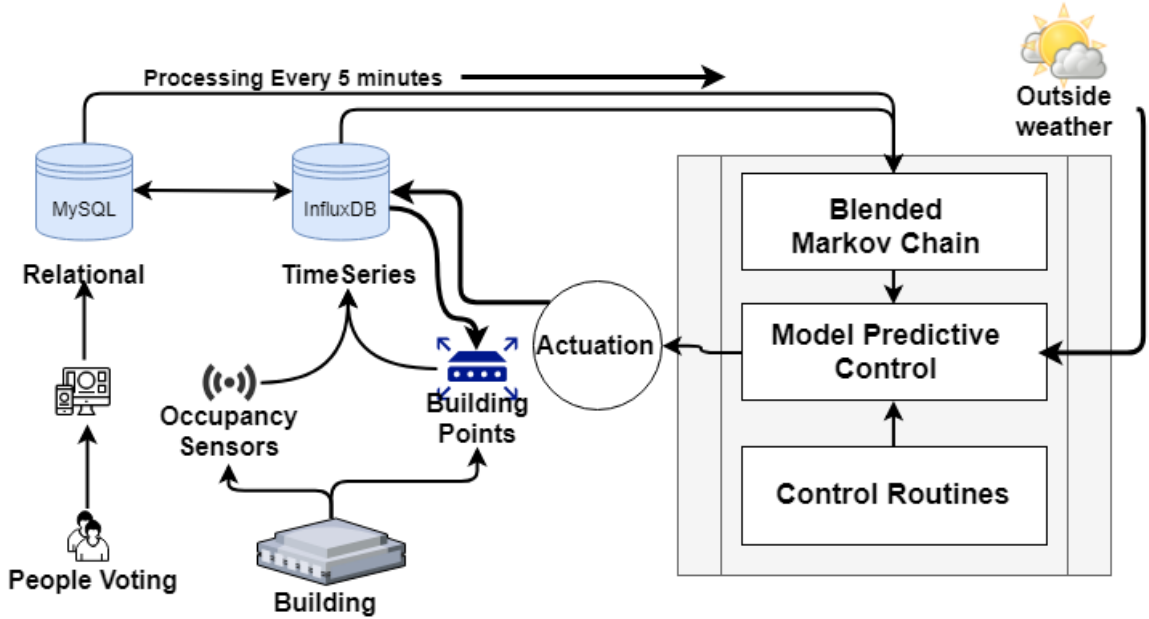


Figure 5.1: OFFICE System Architecture

modules in the processing pipeline add a layer of redundancy on actuation requests to ensure our model predictive control output can never exceed pre-defined minimum and maximum values for control.

The Influx database manages real-time data arriving from periodic data sources. Sensor data is collected by the Building Management System (BMS) periodically throughout the day (every 2 minutes). This sensor data is gathered from sensors such as flow and temperature sensors throughout the HVAC system. While the BMS collects high-resolution data of all points, limitations in the WebCTRL system cause it to severely slow down when third-party applications fetch this data in bulk. As this is a critical system to the facilities management crews on campus, we circumvent this limitation by fetching this data directly from the devices themselves using a data communication protocol for Building Automation and Control Networks called BACnet [LHM12], and stored in the Influx time series database. In addition, the distributed network of occupancy sensors introduced in Section 3.2 periodically sample and transmit occupancy data to a network-enabled border router in the space, where

it is relayed to Influx as well. In contrast, the relational MySQL database manages non-periodic data. For instance, as the thermal comfort application introduced in Section 3.1 receives votes from the users, they are recorded and stored in MySQL to maintain the source of the vote, associated zone for conditioning, etc. To facilitate easy use of this database, the comfort voting application has WEB and Mobile applications built on top of Django(Python) with better support for MySQL.

The time-series database is critical after optimization as well. When the model predictive control finds an optimal solution, it provides a set of setpoints that must be conveyed to the building, and these actuations are all stored in Influx. Once the actuations are validated, a separate module fetches these points from influx and initiates communication with the building. Although the actuations can be set via the Bacnet protocol directly to the actuators in the system, this will cause issues when WebCTRL tries to overwrite the setpoints using its own strategy. Instead, our system communicates with the WebCTRL service's SOAP API as a middleman and overwrites the WebCTRL setpoints, which prevents the interference of the WebCTRL control strategy and allows the facilities techs familiar with the WebCTRL system to have full visibility of the control decisions of the OFFICE system.

As building actuation is done on the scale of minutes, the time it takes to solve for a single optimization problem is less than 5 seconds. Once the optimizer solves the problem, a set of control inputs is obtained, such as the mass flow, discharge temperature, damper positions and supply temperature which is saved in Influx database. The predicted values as explained in section [4] are also saved in the same database. Once everything is in place, we run an actuation script every 5 minutes to read the actuation points from database and these control set points are then transmitted to the BMS using BMS's SOAP based API, which then uses its own built-in PID loops to achieve the set points passed to the building's equipment.

An example of how we control would be setting the point for the mass flow. Our solution contains the flow for each zone which can then be set as required mass flow

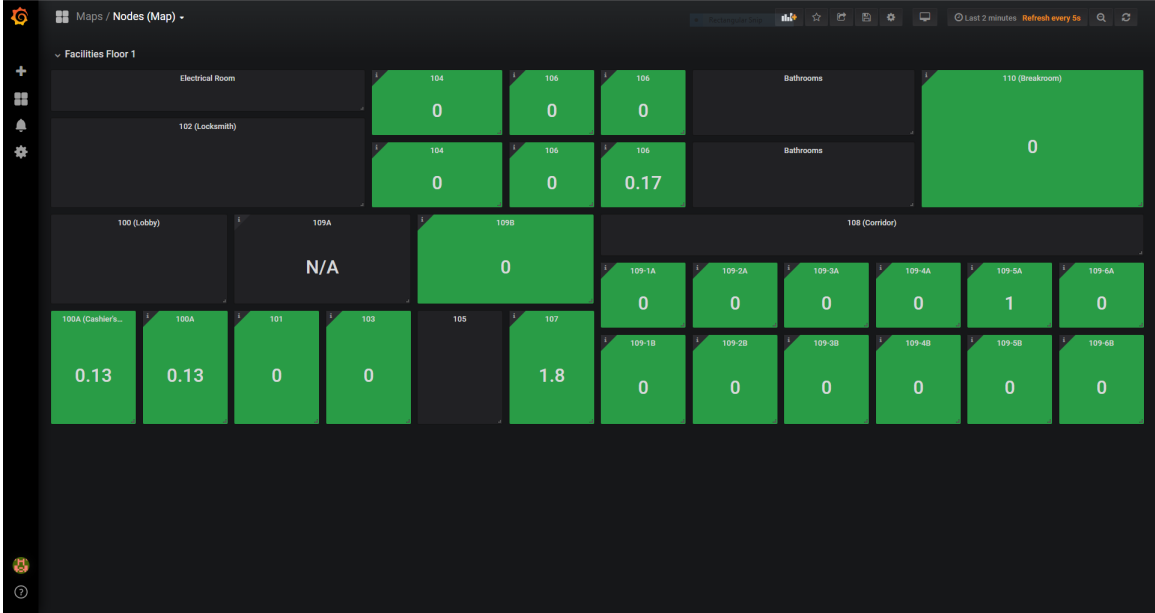


Figure 5.2: OFFICE Grafana Node Map

via the BMS. We only adjust the set-point, and the building is capable of changing the supply fan’s speed and damper position for each zone to meet the needs of all the zones. Every actuation period the points we set include all zones’ mass flow, zones’ supply temperature, discharge temperature at building level. In order to utilize the return air from HVAC system we modify the outside air damper position which help us to conserve energy as the amount of air that needs to be reheated reduces and less energy is spent on reheating the air going in.

Integration with the actual BMS is done in a such way that if the facility management team would like to take control over, it is possible to make an instantly switch from our strategy to the BMS control strategy. As a final layer of oversight, the OFFICE system features monitoring and dashboarding visualization to be used by the facilities techs. We used Grafana [Gra14] and a web application build on NodeJS to integrate a visual monitoring and dashboarding system with our Influx Database to keep track of active wireless sensors for occupancy detection, as well as email alerts in case of a failures, or if a certain zone is not behaving as expected, etc. Figure 5.2



Figure 5.3: OFFICE Zone Temperature Plots

gives an overall picture of the mean number of occupants in a particular room based on a mote and Figure 5.3 provide a historical picture of zone temperature of all the zone. In case a zone temperature goes beyond set bounds, the team gets an email notification with a a visual notification on the dashboard.

In the case that the OFFICE system does something incorrect, we provide the facilities techs with a “kill switch” that immediately terminates OFFICE control and returns to the WebCTRL strategy. To date this feature has not been used.

# CHAPTER 6

## Case Study

### 6.1 Environmental Conditions

To perform a comparison of our system against state-of-the-art building control, it is deployed in a LEED Gold certified building housing the facilities management crew and various administrative staff of our university Figure 6.1. The region within the building under our control is approximately 5000 square feet, and is controlled by 1 air handler unit divided into 8 zones, each provided by a variable air volume (VAV) unit. The HVAC system is controlled by a WebCTRL building management system. In total, the zones cover 4 offices, around 20 cubicles, a conference room, a break room, and an open hallway occupied by two receptionist desks, with each type having its own pattern of usage.

As the university is maintained by the department of facilities and its occupants, our occupants have abnormal hours. Maintenance personnel arrive as early as 5 am and cleaning staff frequently use the building until midnight. In all, there are 18 full-time staff members and several part-time student workers in the building, with more coming in and out as a central operating base. The system was evaluated for a total of 4 weeks ; as the building management system is inactive on weekends, our system analysis includes 20 days of operation on weekdays.




Figure 6.1: 3-D Model of building

## 6.2 Performance Metrics

To determine the merit of our system, we launch it in our building for 20 days, and compare its operation with the baseline WebCTRL strategy with respect to several metrics. In practice, it is infeasible to run the two control strategies at the same time, so it is difficult to compare them under identical conditions; for instance, the outdoor weather changes daily, so the control sequences will change as well. To make sure we have as fair a comparison as possible, for each day of operation under the system, we take the day’s 24-hour outdoor temperature profile, and compare it to our data under WebCTRL operation. When we find the WebCTRL day with the most similar temperature trend, we use it to compare to our system operation. We note that in some cases, more than one day may map to the same WebCTRL day, leading to repeated trends for analysis in WebCTRL data.

The primary purpose of the building management system is to make the indoor environment comfortable for the building occupants. To compare comfort under the two systems, we perform a pre and post-study survey with each occupant in the building to understand how their thermal conditions are under the existing baseline strategy,

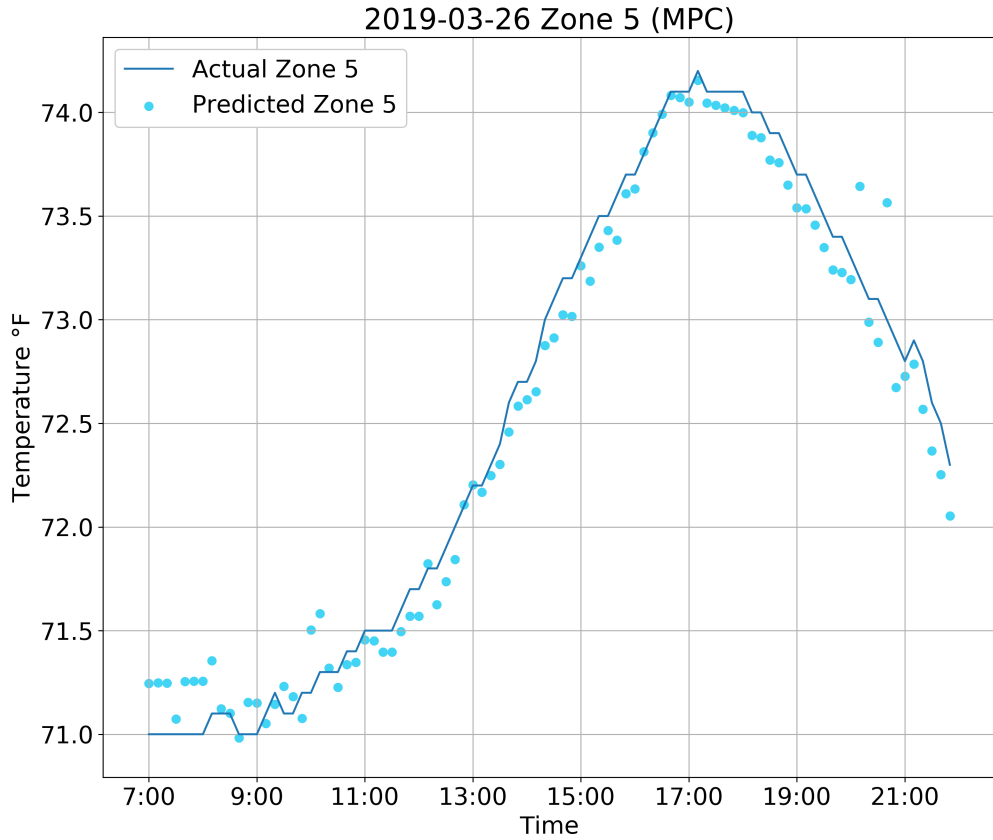


Figure 6.2: Temperature prediction versus ground truth

and how they change under the control of our system. This provides us the clearest view into the comfort provided by the system to *this building's* occupants. In addition, to provide a more quantitative analysis on the more general population, we use Fanger's formula for PMV [Fan70] to estimate comfortable temperature bounds for the "standard" occupant within the current seasonal conditions. With these bounds in mind, we can step through all data from a system's operation, and any time the zone temperature goes outside the bounds, we add the distance from the bound to a sum. By conducting this analysis on both systems, we can directly compare how well in the more general sense the two systems can maintain comfortable conditions.

In addition to comfort, it is critical to the managers of these buildings that the






Figure 6.3: Model error characteristics for each zone

control strategy costs as little money and consumes as little energy as possible. Although minimizing system operational costs and consumed energy are not equivalent goals, we wish to reduce both energy consumption and costs in comparison to the baseline strategies. While we do not have access to direct measurement of energy consumption or costs in the running system, we can use the energy and cost equations derived in Section 4.2 to compute the energy consumed in heating, cooling, and AHU fan using the data trends from the OFFICE system and its most similar weather day under the WebCTRL strategy. While these equations describe the “ideal” energy transfer in the system and will not include energy losses due to heat dissipation or other inevitable sources of loss, it provides us the ability to compare these systems in exactly the same way for comparison.

### 6.3 Model Accuracy

As our system runs its control within the test building, the model predictive control framework is optimizing control actuations based on the model describing the

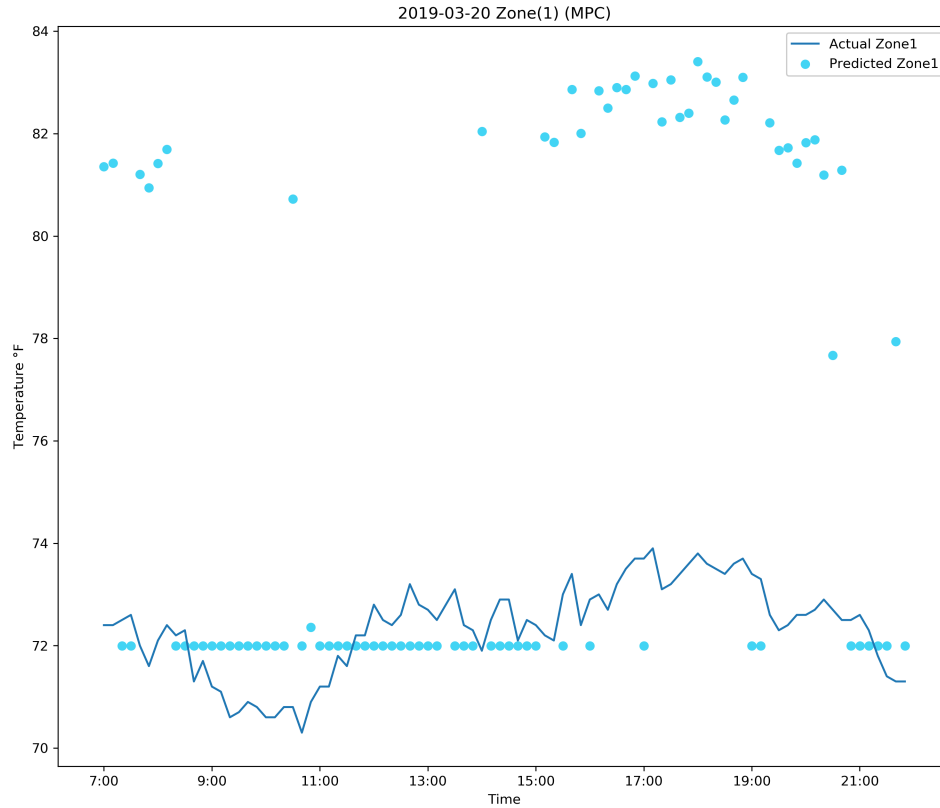


Figure 6.4: Zone 1: Temperature prediction versus ground truth

expected building reaction. To determine whether the thermodynamic model is capturing the effects within the building, we can compare MPC’s predicted temperature at time  $t = +10m$  to the ground truth temperature that actually occurred at time  $t = +10m$ . For instance, Figure 6.2 demonstrates the ground truth zone temperature in Zone 5, against the temperature predicted by the model 10 minutes earlier. We can see in this specific example that the predicted values closely follow the true ones, but that at some times (i.e. 07:00) the model predicts temperatures to increase too quickly, or others (i.e. 19:00) the model predicts the temperature will decrease too quickly.

For a more general and quantitative understanding of the model’s ability to predict

future conditions, the distribution of errors across our 20 days of deployment is shown for each zone in the system in Figure 6.3. In this figure, the center line within each box is the median error for that zone, the upper and lower ends of the boxes define quartiles 1 and 3, the whiskers denote extreme data not considered outliers, and the crosses represent outliers. We can see here that with the exception of zones 1 and 3, the median error for all zones is below 1 degree Fahrenheit, with outliers exceeding 2 degrees. While we believe this is reasonable performance in the more general case, it is clear that the model had difficulty making accurate predictions for zones 1 and 3. Zone 1 is a break room in the building, and has many characteristics that make prediction difficult. As a break room, occupancy patterns were found to be very sporadic and appliances (1 fridge, 2 vending machines, 1 printer) cause irregular thermal load in the room that was difficult to model. Zone 1's error in prediction can be seen in Figure 6.4. Zone 3, an open cubicle area, had the main door to the outside, which resulted in a large influx of outside air every time somebody entered or exited the building. It is important to note that while some zones have higher modeling error, the OFFICE framework will be able to provide control as long as the model predicts future states in the direction of the ground truth. In these cases, OFFICE's periodic re-solve of control optimization with corrected initial conditions will guide the solution towards the desired state, even if it is sub-optimal w.r.t. cost. To demonstrate this point, Section 7 will show that the OFFICE system is capable of improved occupant comfort *and* energy savings despite this challenge in modelling.

# CHAPTER 7

## Results

For our analysis, we want to compare the state-of-art building control that is operational in the top of the line green buildings with our OFFICE control strategy. We think this is valid baseline comparison as this represents the best strategy that is implemented in current top quality building stock in the U.S. As mentioned in Section 6.2, we are interested in comparing both the quality of thermal comfort and the energy use/cost results between the two systems.

### 7.1 Thermal Comfort Analysis

Table 7.1 shows the results from the pre-survey and the post-survey. The first thing to notice is that the satisfaction level in the pre-survey for the WebCTRL baseline strategy is very good. Only 25% of users are dissatisfied, with 50% of all users being satisfied or somewhat satisfied. Previous work has shown that the percentage of users

Satisfaction level	Pre-survey	Post-survey
Satisfied or very satisfied	25.0%	30.0%
Somewhat satisfied	25.0%	50.0%
Neutral	25.0%	20.0%
Somewhat dissatisfied	18.75 %	<b>0.0%</b>
Dissatisfied	6.25%	<b>0.0%</b>

Table 7.1: Room thermal satisfaction survey report

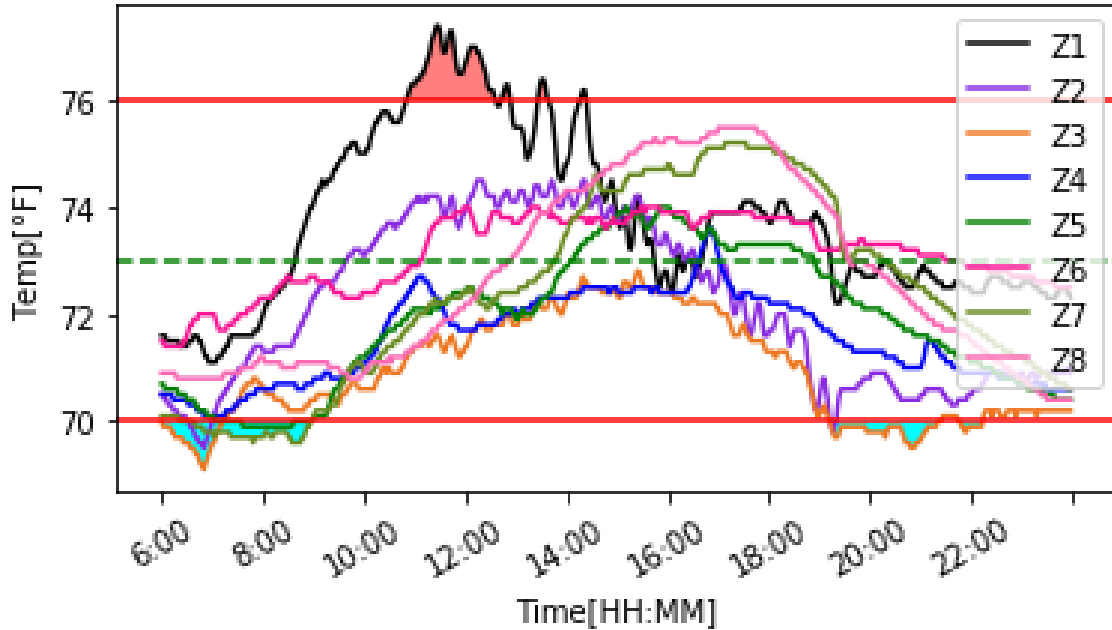


Figure 7.1: Example of thermal comfort for all zones in WebCTRL

	Total Sum	No of Days	Average
WebCTRL	3543.3	84	42.2
OFFICE	459.8	20	23.0

Table 7.2: Quality of Thermal Comfort

dissatisfied was in general much larger, varying from 50% to 75% [EC12, WBE16]. This is not completely surprising though, as the building went through LEED Gold certification process and two very experienced HVAC engineer tuned up the control routines in the building to make them as efficient as possible, debugging comfort issues with users in the process. In the post-survey column, we see the results of the post-survey after running the OFFICE system. The level of dissatisfaction is zero percent, and 80% of the users were satisfied or somewhat satisfied, with the remaining 20% being neutral.

We want to take a closer look as to why WebCTRL has significantly worse quality

of thermal comfort than OFFICE. Figure 7.1 shows an example of thermal comfort for all the zones when running WebCtrl for a single day. As mentioned in Section 6.2, we use Fanger’s formula for the Predictive Mean Vote (PMV) [Fan70] to estimate comfortable temperature bounds for the “standard” occupant within the current seasonal conditions, as defined by ASHRAE standard 55 [ASH07a]. Office-like environments do not require special thermal needs, so they are considered Class C environments according to the classification scheme included in ISO 7730 [Iso05]. The maximum high/low end of the comfort range for Class C environments has PMV values of  $\pm 0.7$ , which we determined to be a temperature range between 70°F and 76°F. From Figure 7.1 we see that over the course of the day, many zones get close to the the temperature limits, and specifically zones 1 and 3 exceeds the low and high limits several times during the day. The figure shows clearly that WebCTRL exceeds the comfort temperature bounds and would tend to provide poor quality of thermal comfort to its occupants. In order to provide more quantitative results over the course of many more days, we step through all data from both WebCTRL and OFFICE, and any time the zone temperature goes outside the bounds, we add the distance from the bound to sum, i.e. we calculate the L1 distance when bounds are exceeded. The results are shown in Table 7.3. From the table, we see that WebCTRL produces a significantly larger distance when comfort bounds are violated, with an average distance across all zones of 42.2/day versus OFFICE 23.0/day, an 83.5% decrease in quality of comfort by WebCTRL compared to OFFICE.

Figure 7.2 shows the thermal preferences vote distribution for all the users in our building. We see that  $\sim 41\%$  of all the users’ votes were on the hot/warm side of the scale, with  $\sim 32\%$  of all the users’ votes on the cold/cool side, and  $\sim 27\%$  of all users’ vote being neutral. This figure provides an indication that temperatures might be on the warmer side for the majority of the zones. To verify this intuition based on the users thermal preferences, we plot in Figure 7.3 the zone average temperature variation from all zones across different times of the day for a sample experimental day.

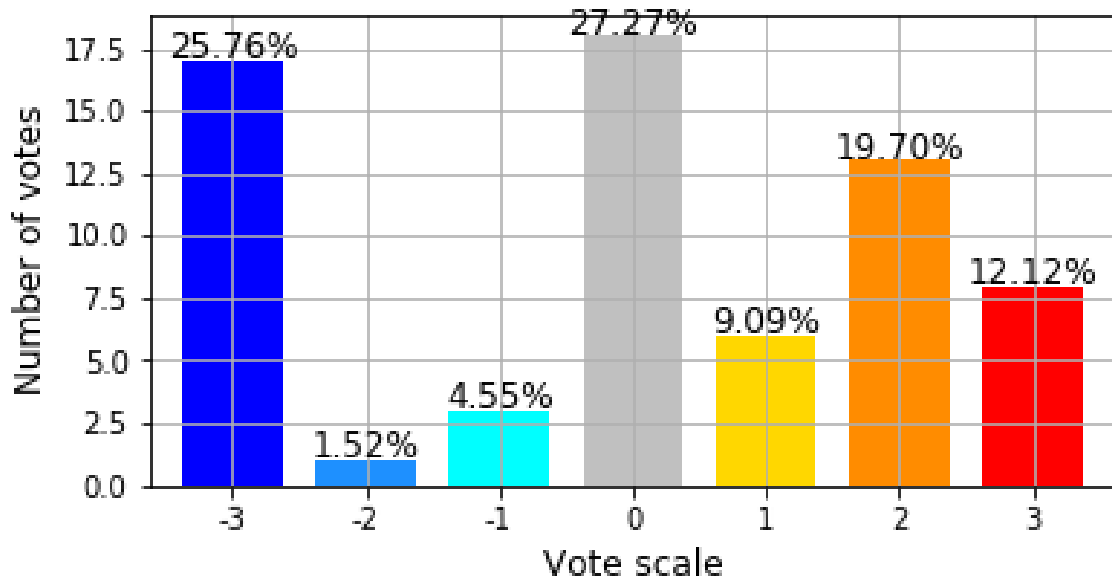


Figure 7.2: Thermal preferences vote distribution for the users in our building. Values [-3,-2,-1,0,1,2,3] corresponds to [Cold, Cool, Slightly Cool, Neutral, Slightly Warm, Warm, Hot] in our comfort app.




Figure 7.3: Zone temperature variation across the day for all zones for a selected day

	Average Cost(USD)	Average Energy(Kwh)
WebCTRL	9.235	112.106487
OFFICE	8.242	107.161204

Table 7.3: Savings

We see from the figure that the majority of the zones experienced warm temperatures, except in the mornings and except zone 3 across most of the day. Since voting patterns may be changing across the day depending on the zone temperature, we plot the voting data for cold, warm and neutral values as a function of the time of the day in Figure 7.4. The data was binned at the beginning of the time slot, so a vote indicated at noon could have occurred at anytime since 12:00pm to 12:59pm and so forth for the rest of the hours during the day. We see that there is a clear pattern from the figure where most of the cold votes occur in the morning across all days, while the majority of the warm votes occur in the afternoon. A peculiar thing we noticed is that a large concentration of votes occur between 12:00 and 12:59pm. This time also coincides with the lunch break for the majority of users. We hypothesize that this may be the effect of the temperature transitions that occur in the majority of the zones, from the colder mornings to the warmer afternoon, and the transitions may motivate users to vote more to protect their “thermal well-being”. Another explanation might be that users may have more time during the lunch break for tasks other than work, and they may pay more attention to thermal comfort and use the comfort app more often. However, more investigation is required as to the reasons why this happens.

## 7.2 Cost and Energy Analysis

Our other primary metrics of success are the cost and energy profiles of the running system. As discussed in Chapter 4, the optimization problem will minimize the cost of operation, but as the occupants of the building have the opportunity to influence



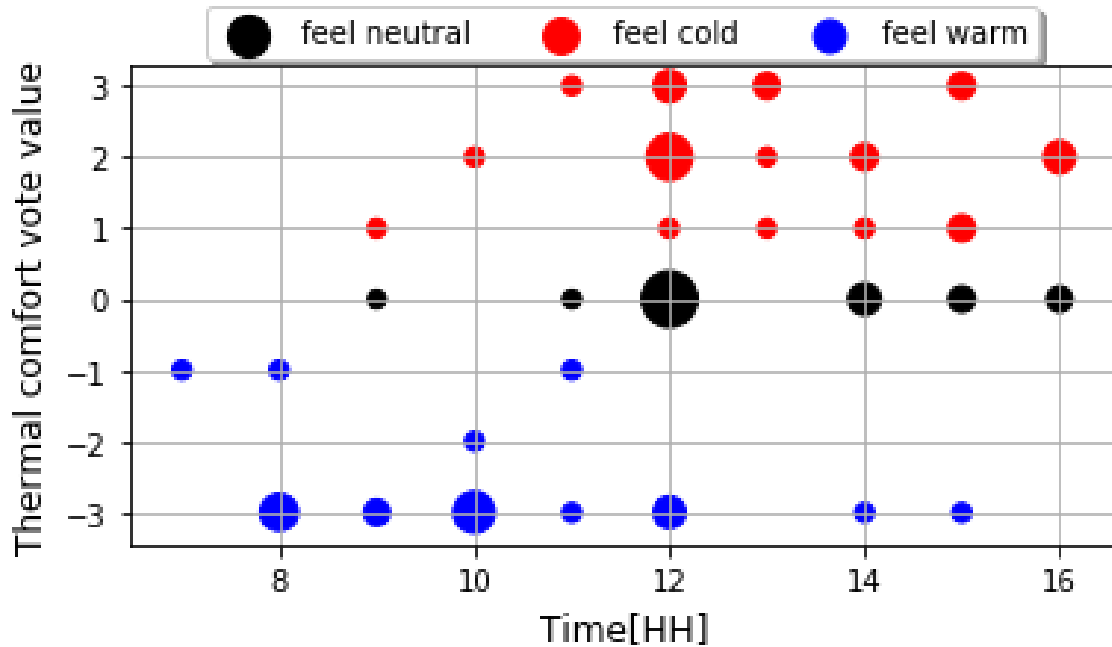


Figure 7.4: Votes distribution in our comfort app as a function of the time of the day.

the control routines by voting for comfort in their space, the OFFICE system may have more restrictive temperature setpoints, potentially increasing these metrics. As discussed in Chapter 6, to perform a comparison of the OFFICE and WebCTRL control strategies, each day of OFFICE operation is matched to the closest day of WebCTRL operation with respect to hourly outdoor weather trends, allowing us to compare the two systems, even when they aren't operational at the same time.

The energy consumption of the two buildings is compared in Figure 7.6, where the solid color bars represent the energy profile under the OFFICE system, and the corresponding hatched bar represents the energy profile of WebCTRL's most similar day. In addition, each bar is broken into the separate heating and cooling costs, stacked on top of each other. We can see here that a bit more energy is spent in cooling the air, which is expected due to the slightly warmer outdoor temperature of the spring season. It may be counter-intuitive that both heating and cooling is occurring in this system, but this is expected of a terminal reheat system such as




Figure 7.5: Operation cost of OFFICE vs WebCTRL

this; the cooling is done at the AHU level, bringing the air to a cool temperature to be distributed to the building, and then terminal reheat at the VAV level warms the air to the temperature required for the space. In our system in particular, each zone may have differing discharge air requirements due to occupant preference and voting, so this terminal reheat is common to meet these demands.

### 7.2.1 Cost Analysis —& Energy Analysis

Figure 7.5 shows the cost comparison between the OFFICE and WebCTRL systems across the 20 days of deployment. The reason Figures 7.6 and 7.5 are not directly proportional is due to the differences in energy cost from gas and electrical sources. As electricity costs approximately 5x more than gas, the majority of costs are due to electrical usage in building cooling. In terms of cost, the OFFICE system required 8.24\$ per day vs WebCTRL’s 9.23\$ per day, a 10.8% overall savings.

As shown in Figure 7.6, the OFFICE system consumed an average 107.16kWh per day vs WebCTRL’s 112.11kWh per day, a 4.4% savings in energy consumption under the OFFICE system.




Figure 7.6: Energy consumption of OFFICE vs WebCTRL

# CHAPTER 8

## Discussion

In this section, we discuss some of our experiences and observations while controlling a building’s HVAC system using our OFFICE framework.

There are several limitations to our systems that were essentially discovered only after extensive experimentation. The first issue is that buildings have faulty sensor and equipment, and no control strategy can replace good maintenance and fault detection and diagnostic of the building. For example, in very specific zones we observed days when the zone discharge temperature was 10°F higher than the zone temperature, but the zone failed to heat up. Misaligned sensors make our thermodynamic model appear incorrect, when in fact the issue is faulty hardware. Another example was the change in the connection polarity of an actuator in the economizer. When instructing the system to use 100% outside air, the system went into full recirculation with not outside air and vice-versa. We were able to debug the system in the initial trial runs and we fixed the problem in our software framework by simply inverting the logic without changing the hardware, but on large systems developers should be aware of these type of problems. Another limitation of the system is in our gray-box approach for the thermodynamic model. We decided to use a simple enough model that could be used in optimization control loop, but powerful enough to get the “derivative right” even at the expense of some amplitude error in some corner cases. In Section 6.3 we explored in detail our model accuracy, explaining some of the reasons why we have larger model errors for zone 1 (breakout room with appliances) and zone 3 (open area with cubicles exposed to a door being frequently

opened). We think there is room for improvement here, and perhaps we should explore thermodynamic models that are able to cope better with variable loads and wall (door) resistant parameters that expect certain levels of dynamics.

In addition to this, we note that MPC makes the assumption that the changing of flow and temperature setpoints in the system occur very fast relative to our actuation frequency of 10 minutes. In running the system, however, we found that adjustment to temperatures and flows at the AHU level often took several minutes, which will cause the model's prediction to be inaccurate immediately after changes are made to the setpoints as the changes are taking effect.

Another point worth mentioning is that the quality of the weather forecasting data is highly dependent on the station location. During the experiment we compared outside temperature using a temperature thermometer with the weather data from the forecasting API [wun19]. Although the difference was not significant, having the actual outside temperature the building is exposed to, would increase the accuracy of the optimization results. For future work, we could consider deploying our own sensor near the outside input vent of the main AHU loop.

# CHAPTER 9

## Conclusion

In this thesis, we covered several topics with a goal to balance the tradeoff between energy cost and user comfort in the operation of HVAC systems in commercial buildings. We collect missing information from current HVAC systems, namely

- a) real-time and predictive occupancy using wireless sensor nodes,
- b) human-in-the-loop comfort data gathered through mobile comfort apps,
- c) develop a data-driven gray-box model where the parameters of the model are learned from the building operational data.

Together with weather forecast information, this data is fed into OFFICE to minimize the energy cost while satisfying user comfort constraints. We tested our system for over 4 weeks in a LEED Gold certified building, which was recently optimized by two experienced HVAC engineers for optimal operation, where 20+ users in 8 zones do their daily tasks. We showed a significant reduction in dissatisfaction, going from 25% with the baseline strategy to 0% dissatisfaction with OFFICE. Despite this improvement in quality of thermal comfort, we were able to also show a decrease of the energy cost by more than 10% compared with the baseline strategy.

There are many avenues to be taken for future work on the basis of this work. The optimization problem developed in the end can be further expanded to take additional elements into consideration such as non-constant thermal load. As discussed in Chapter 8, extra sensors can be deployed to get accurate measurement of weather data to make the system more efficient.

## REFERENCES

- [ABD11] Yuvraj Agarwal, Bharathan Balaji, Seemanta Dutta, Rajesh Gupta, and Thomas Weng. “Duty-Cycling Buildings Aggressively: The Next Frontier in HVAC Control.” In *IPSN*, 2011.
- [Adm] US Energy Information Administration. “How much energy is consumed in U.S. residential and commercial buildings?” <https://www.eia.gov/tools/faqs/faq.php?id=86&t=1>.
- [ASH07a] “ASHRAE Standard 55: Thermal Environmental Conditions for Human Occupancy.” American Society of Heating, Refrigeration and Air-Conditioning Engineers, Inc., 2007.
- [ASH07b] “ASHRAE Standard 62.1: Ventilation for Acceptable Indoor Air Quality.” American Society of Heating, Refrigeration and Air-Conditioning Engineers, Inc., 2007.
- [BC14] Alex Beltran and Alberto E Cerpa. “Optimal HVAC building control with occupancy prediction.” In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, pp. 168–171. ACM, ACM, 2014.
- [BEC13] Alex Beltran, Varick L Erickson, and Alberto E Cerpa. “Thermosense: Occupancy thermal based sensing for hvac control.” In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, pp. 1–8. ACM, ACM, 2013.
- [BEK14] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. “Julia: A Fresh Approach to Numerical Computing.” *CoRR*, 2014.
- [BLT10] Athanasios Bami, Dimitrios Lymberopoulos, Thiago Teixeira, and Andreas Savvides. “The BehaviorScope framework for enabling ambient assisted living.” *Personal Ubiquitous Comput.*, **14**(6):473–487, September 2010.
- [BTG13] Bharathan Balaji, Hidetoshi Teraoka, Rajesh Gupta, and Yuvraj Agarwal. “ZonePAC: Zonal Power Estimation and Control via HVAC Metering and Occupant Feedback.” In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys’13, pp. 18:1–18:8, New York, NY, USA, 2013. ACM.
- [Bur14] “Human-Building Interaction Framework for Personalized Thermal Comfort-Driven Systems in Office Buildings.” *Journal of Computing in Civil Engineering*, **28**(1):2–16, 2014.

- [CAN14] M Castilla, JD Álvarez, JE Normey-Rico, and F Rodríguez. “Thermal comfort control using a non-linear MPC strategy: A real case of study in a bioclimatic building.” *Journal of Process Control*, **24**(6):703–713, 2014.
- [Cou17] The U.S. Green Building Council. “LEED.”, 2017.
- [D12] Ltd. D&R International. *2011 Building Energy Data Book*. U.S. Dept. of Energy, 2012.
- [EAC] V. L. Erickson, S. Achleitner, and A. E. Cerpa. “Poem: Power-efficient Occupancy-based Energy Management System.” In *IPSN’13*.
- [EAC13] Varick L Erickson, Stefan Achleitner, and Alberto E Cerpa. “POEM: Power-efficient occupancy-based energy management system.” In *Proceedings of the 12th international conference on Information processing in sensor networks*, pp. 203–216. ACM, ACM, 2013.
- [EC] Varick Erickson and Alberto 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 (BuildSys 2012)*.
- [EC12] Varick L Erickson and Alberto 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*, pp. 9–16. ACM, 2012.
- [ECC11] Varick L Erickson, Miguel Á Carreira-Perpiñán, and Alberto E Cerpa. “OBSERVE: Occupancy-based system for efficient reduction of HVAC energy.” In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pp. 258–269. IEEE, 2011.
- [ECC14] Varick L Erickson, Miguel Á Carreira-Perpiñán, and Alberto E Cerpa. “Occupancy modeling and prediction for building energy management.” *ACM Transactions on Sensor Networks (TOSN)*, **10**(3):42, 2014.
- [Fan70] Poul O Fanger et al. “Thermal comfort. Analysis and applications in environmental engineering.” *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.
- [FFS06] William Fisk, David Faulkner, and Douglas Sullivan. “Accuracy of CO2 sensors in commercial buildings: a pilot study.” Technical report, LBNL, 2006.
- [GIB12] Siddharth Goyal, Herbert A Ingley, and Prabir Barooah. “Zone-level control algorithms based on occupancy information for energy efficient buildings.” In *2012 American Control Conference (ACC)*, pp. 3063–3068. IEEE, 2012.



- [Gra14] “Grafana.”, 2014.
- [HW13] Lam Abraham Hang-yat and Dan 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*, BuildSys’13, pp. 21:1–21:8, New York, NY, USA, 2013. ACM.
- [inf13] “InfluxData.”, 2013.
- [Iso05] En Iso. “7730: 2005.” *Ergonomics of the thermal environment-Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria*, 2005.
- [JB12] Farrokh Jazizadeh and Burcin 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*, BuildSys ’12, pp. 1–8, New York, NY, USA, 2012. ACM.
- [JGB13] Farrokh Jazizadeh, Ali Ghahramani, Burcin Becerik-Gerber, Tatiana Kichkaylo, and Michael Orosz. “Human-building interaction framework for personalized thermal comfort-driven systems in office buildings.” *Journal of Computing in Civil Engineering*, **28**(1):2–16, 2013.
- [KB11] Anthony Kelman and Francesco Borrelli. “Bilinear model predictive control of a HVAC system using sequential quadratic programming.” *IFAC Proceedings Volumes*, **44**(1):9869–9874, 2011.
- [KJD09] Ankur Kamthe, Lun Jiang, Matt Dudys, and Alberto Cerpa. “SCOPES: Smart Cameras Object Position Estimation System.” In *EWSN*, 2009.
- [KMB13a] Anthony Kelman, Yudong Ma, and Francesco Borrelli. “Analysis of local optima in predictive control for energy efficient buildings.” *Journal of Building Performance Simulation*, **6**(3):236–255, 2013.
- [KMB13b] Anthony Kelman, Yudong Ma, and Francesco Borrelli. “Analysis of local optima in predictive control for energy efficient buildings.” *Journal of Building Performance Simulation*, **6**(3):236–255, 2013.
- [KNO01] Neil E Klepeis, William C Nelson, Wayne R Ott, John P Robinson, Andy M Tsang, Paul Switzer, Joseph V Behar, Stephen C Hern, and William H Engelmann. “The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants.” *Journal of Exposure Science and Environmental Epidemiology*, **11**(3):231, 2001.

- [LHD09] Khee Poh Lam, Michael Hoyneck, Bing Dong, Burton Andrews, Yun-Shang Chiou, Rui Chang, Diego Benitez, and Joonho Choi. “Occupancy Detection Through an Extensive Environmental Sensor Network in an Open-Plan Office Building.” In *International Building Performance Simulation Association*, 2009.
- [LHM12] S Liaisons, R Hall, M Modera, C Neilson, B Isler, M Osborne, DP Alexander, CL Brumley, CH Copass, SE Dinges, et al. “BACnet-A Data Communication Protocol for Building Automation and Control Networks.” *ANSI/ASHRAE Standard*, **135**, 2012.
- [LSS10] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. “The smart thermostat: using occupancy sensors to save energy in homes.” In *SenSys*, 2010.
- [MB12] Yudong Ma and Francesco Borrelli. “Fast stochastic predictive control for building temperature regulation.” In *2012 American Control Conference (ACC)*, pp. 3075–3080. IEEE, IEEE, 2012.
- [MCM12] Sunil Mamidi, Yu-Han Chang, and Rajiv Maheswaran. “Improving building energy efficiency with a network of sensing, learning and prediction agents.” In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 45–52. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [MJK13] Parisa Mansourifard, Farrokh Jazizadeh, Bhaskar Krishnamachari, and Burcin 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*, BuildSys’13, pp. 20:1–20:8, New York, NY, USA, 2013. ACM.
- [MKD12] Yudong Ma, Anthony Kelman, Allan Daly, and Francesco Borrelli. “Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments.” *IEEE control systems magazine*, **32**(1):44–64, 2012.
- [MMB15] Yudong Ma, Jadranko Matusko, and Francesco Borrelli. “Stochastic model predictive control for building HVAC systems: Complexity and conservatism.” *IEEE Transactions on Control Systems Technology*, **23**(1):101–116, 2015.
- [MMC11] S. Mamidi, R. Maheswaran, and Y. Chang. “Smart Sensing, Estimation, and Prediction for Efficient Building Energy Management.” In *Multi-agent Smart Computing Workshop*, 2011.
- [MyS95] “MySQLDB.”, 1995.

- [ole73] “American Society of Heating, Refrigerating and Air-Conditioning Engineers, ASHRAE Standard : Standards for Natural and Mechanical Ventilation. New York :The Society,.” 1973.
- [PKJ13] Sean Purdon, Branislav Kusy, Raja Jurdak, and Geoffrey Challen. “Model-free HVAC control using occupant feedback.” In *38th Annual IEEE Conference on Local Computer Networks-Workshops*, pp. 84–92. IEEE, IEEE, 2013.
- [PMV13] Alessandra Parisio, Marco Molinari, Damiano Varagnolo, and Karl Henrik Johansson. “A scenario-based predictive control approach to building HVAC management systems.” In *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, pp. 428–435. IEEE, 2013.
- [PVR13] Alessandra Parisio, Damiano Varagnolo, Daniel Risberg, Giorgio Pattarello, Marco Molinari, and Karl H Johansson. “Randomized model predictive control for HVAC systems.” In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, pp. 1–8. ACM, 2013.
- [Ram69] J. B. Ramsey. “Tests for Specification Errors in Classical Linear Least-Squares Regression Analysis.” *Royal Statistical Society*, 1969.
- [RES10] Sasank Reddy, Deborah Estrin, and Mani Srivastava. “Recruitment Framework for Participatory Sensing Data Collections.” In *Proceedings of the 8th International Conference on Pervasive Computing, Pervasive’10*, pp. 138–155, Berlin, Heidelberg, 2010. Springer-Verlag.
- [Rob14] Building Robotics. “Case Study: Johnson Controls Inc.” Technical report, 300 Frank H Ogawa Plaza #614 — Oakland, CA, January 2014.
- [RVL15] Ricardo Forgiarini Rupp, Natalia Giraldo Vásquez, and Roberto Lamberts. “A review of human thermal comfort in the built environment.” *Energy and Buildings*, **105**:178–205, 2015.
- [SAS17] Fisayo Caleb Sangogboye, Krzysztof Arendt, Ashok Singh, Christian T Veje, Mikkel Baun Kjærgaard, and Bo Nørregaard Jørgensen. “Performance comparison of occupancy count estimation and prediction with common versus dedicated sensors for building model predictive control.” In *Building Simulation*, volume 10, pp. 829–843. Springer, Spriner, 2017.
- [SBK11] James Scott, A.J. Bernheim Brush, John Krumm, Brian Meyers, Mike Hazas, Steve Hodges, and Nicolas Villar. “PreHeat: Controlling Home Heating Using Occupancy Prediction.” In *Proceedings of UbiComp 2011*. ACM, September 2011.

- [SFC18] Gianluca Serale, Massimo Fiorentini, Alfonso Capozzoli, Daniele Bernardini, and Alberto Bemporad. “Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities.” *Energies*, **11**(3):631, 2018.
- [SGM16] David Sturzenegger, Dimitrios Gyalistras, Manfred Morari, and Roy S Smith. “Model predictive climate control of a swiss office building: Implementation, results, and cost–benefit analysis.” *IEEE Transactions on Control Systems Technology*, **24**(1):1–12, 2016.
- [WBE16] Daniel A Winkler, Alex Beltran, Niloufar P Esfahani, Paul P Maglio, and Alberto E Cerpa. “FORCES: feedback and control for occupants to refine comfort and energy savings.” In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 1188–1199. ACM, ACM, 2016.
- [Wik19] “WikiIoT.”, 2019.
- [WPL19] Zhe Wang, Thomas Parkinson, Peixian Li, Borong Lin, and Tianzhen Hong. “The Squeaky wheel: Machine learning for anomaly detection in subjective thermal comfort votes.” *Building and Environment*, 2019.
- [wun19] “Wunderground.”, 2019.
- [YAL14] Ray Yun, Azizan Aziz, Bertrand Lasternas, Chenlu Zhang, Vivian Loftness, Peter Scupelli, Yunjeong Mo, Jie Zhao, and Nana Wilberforce. “The design and evaluation of intelligent energy dashboard for sustainability in the workplace.” In *International Conference of Design, User Experience, and Usability*, pp. 605–615. Springer International Publishing, 2014.