Technical University of Crete

School of Electrical and Computer Engineering

Balancing expected thermal discomfort and Heating-Ventilation-Air Conditioning operation cost in buildings with dynamic occupancy schedules

Author:

Michail Katsigiannis

Committee:

Georgios Chalkiadakis

Konstantinos Kalaitzakis

Athanasios Aris Panagopoulos

A thesis submitted in fulfillment of the requirements for the degree of

Diploma in Electrical and Computer Engineering

July 2020

Abstract

Optimizing the heating, ventilation and air condition (HVAC) process considers a multi-objective optimization task with the main objectives being energy cost and thermal discomfort. Typically, when the occupancy schedule of a building is known in advance, a thermostat can be set to predetermined values in order to minimize thermal discomfort. However, in dynamic occupancy schedule scenarios, ensuring minimum thermal discomfort at all times is highly energy inefficient. In more detail, even with the most accurate occupancy prediction algorithms some uncertainty on the predicted schedule is retained. Minimizing thermal discomfort, even in the slightest chance of occupancy, introduces unnecessary and unacceptable cost. In this context, our goal in this thesis was to investigate HVAC control optimization algorithms that incorporate occupancy predictions in a feasible manner. In more detail, we compare two algorithms derived from the literature against a novel algorithm that we propose here, with respect to applicability, effectiveness, efficiency and usability criteria. Our comparison shows that each one of the algorithms possesses certain advantages and disadvantages with respect to the abovementioned criteria. Our approach performs similar to state of the art approaches and offers increased usability since it relies on a single intuitive parameter.

Περίληψη

Η επίτευξη της βέλτιστης θέρμανσης, εξαερισμού και κλιματισμού (HVAC) κτιρίων είναι μία πολύπλοκη διαδικασία που προαπαιτεί μια εργασία μαθηματικής βελτιστοποίησης πολλαπλών στόχων, οι κυριότεροι εκ των οποίων είναι το κόστος ενέργειας και τη θερμική δυσφορία. Συνήθως, όταν σε ένα κτίριο είναι γνωστό το χρονοδιάγραμμα χρήσης του κτιρίου εκ των προτέρων, τότε ένας θερμοστάτης μπορεί να ρυθμιστεί σε προκαθορισμένες τιμές, προκειμένου να ελαχιστοποιηθεί η θερμική δυσφορία. Ωστόσο, σε σενάρια που υπάρχει δυναμικό χρονοδιάγραμμα χρήσης, η εξασφάλιση της ελάχιστης θερμικής δυσφορίας ανά πάσα στιγμή είναι εξαιρετικά ενεργειακά αναποτελεσματική. Πιο συγκεκριμένα, ακόμα και με τους πιο ακριβείς αλγόριθμους πρόβλεψης πληρότητας των χώρων ενός κτιρίου, διατηρείται κάποια αβεβαιότητα σχετικά με το προβλεπόμενο χρονοδιάγραμμα χρήσης. Η ελαχιστοποίηση της θερμικής δυσφορίας, ακόμη και στις ελάχιστες πιθανότητες πληρότητας, εισάγει περιττό και μη αποδεκτό κόστος. Σε αυτό το πλαίσιο, ο στόχος μας είναι να διερευνήσουμε αλγορίθμους βελτιστοποίησης ελέγχου συστημάτων HVAC που ενσωματώνουν προβλέψεις πληρότητας με εφικτό τρόπο. Συγκεκριμένα, στην εργασία μας συγκρίνουμε δύο αλγόριθμους που προέρχονται από τη βιβλιογραφία, με έναν νέο αλγόριθμο που προτείνουμε εδώ, λαμβάνοντας υπ΄ όψιν στη σύγκριση κριτήρια εφαρμογής, αποτελεσματικότητας, αποδοτικότητας και χρηστικότητας. Η πειραματική μας σύγκρισή υποδεικνύει ότι καθένας από τους αλγορίθμους διαθέτει ορισμένα πλεονεκτήματα και μειονεκτήματα σε σχέση με τα προαναφερθέντα κριτήρια. Η προσέγγισή μας παρέχει παρόμοια απόδοση με τις σύγχρονες προσεγγίσεις και προσφέρει αυξημένη χρηστικότητα, καθώς βασίζεται σε μόνο μία παράμετρο που έχει διαισθητική ερμηνεία.

Contents

List of Figures

List of Symbols and Abbreviations

λ	Weighting parameter.
a	HVAC control action vector.
c	Regression coefficient.
$Disc$	Instantaneous thermal discomfort.
$E[Disc]$	Expected instantaneous thermal discomfort.
i	Additional variables vector.
J	Weighted sum.
O	Probabilistic Occupancy Estimates.
t	Time horizon.
T^{down}	Lower expanded comfort band limit.
T^{down}	Lower limit of the comfort band.
T^{in}	Indoor temperature.
T^{maxup}	Upper absolute comfort band limit.
$T^{mindown}$	Lower absolute comfort band limit.
T^{up}	Upper expanded comfort band limit.
T^{up}	Upper limit of the comfort band.

TM	Thermal model.
x	Thermal state vector.
DAG	Directed acyclic graph.
DSM	Demand side management.
HVAC	Heating, ventilation and air condition.
IOT	Internet of things.
MPC	Model predictive control.
RMSE	Root-mean square error.

Acknowledgements

I would first like to express my deep gratitude towards my supervisor, Dr. George Chalkiadakis, who inspired me to to conduct research in this field, and entrusted me with the opportunity to handle the responsibilities required for the completion of this research.

Secondly, i would like to express my deep gratitude for Dr. Athanasios Aris Panagopoulos for his advice and assistance. His valuable and constructive help during the planning and development of this research work, his offer of all the resources that were critical during the research and his willingness to give his time so generously have been very much appreciated.

I would also like to acknowledge all of my friends and classmates, that provided me with their support and encouragement during the time of my studies.

Finally, I would also like to thank my family that provided me with all the support that I needed and for creating for me a loving and encouraging environment, in which I was able to finish the entirety of my studies. They are the ones that kept me going on and none of this work would have been possible without them.

Chapter 1

Introduction

One of the biggest challenges of the electrical grid (and of all electrical systems) operation is to match supply and demand. Any lack in matching supply and demand might lead to blackouts or overheating of the electrical grid. Traditionally, this challenge has been tackled by having in operation highly dispatchable supply in order to match a fluctuating and largely unpredictable electricity demand. To this end, fossil fuels provided a great energy resource that is highly dispatchable and is able to support such an operation. In more detail, fossil fuels based generators are very fast in increasing or decreasing the energy generation according to the needs of the grid operator. That said, in recent years we are experiencing a shift in the paradigm of the electricity grid operation, both in the production and the consumption of energy; a rapid re-engineering of the electrical grid. Chief among them are the electrification of heating and transportation, along with the increasing introduction of renewable energy generators, such as wind turbines and photovoltaic systems in the energy generation mix.

This rapid re-engineering of the electrical grid and the electrical grid operation serves two main goals: (i) the support of an ever increasing energy demand, and (ii) the reduction of our reliance on fossil fuels. In more detail, in recent years we rely on electrical devices, or energy consuming devices, on almost all every day activities, ranging from entertainment and healthcare, to security and transportation. This, in turn explodes our energy requirements. In this context, the electrification of heating and transportation can aid, given the fact that electricity based engines and electricity based heating generators (such as heat pump based technology) are more energy efficient compared to fossil-fueled-based ones.

Moreover, our reliance on fossil fuel has been accused as the main reason for the climate change. Such as, decreasing this reliance on fossil fuel can aid in countering these changes. Also, reducing our reliance on fossil fuel can also support countries and regions that do not have fossil fuel reserves. In this context, the electrification of heating and transportation is aiding the reduction of this reliance, compared to internal combustion engines and fossil fuel based heating. In addition, the introduction of renewable energy generation into the energy generation mix can ensure that the electricity provided to these electricity based technologies comes from renewable generation and not fossil fuel.

That said, the ever-increasing penetration of renewable energy resources into the energy generation mix raises stability issues for the electrical grid. In more detail, the energy generated by such systems is not human controlled, but rather weather driven, as it relies on the prevailing weather conditions. The energy generated by photovoltaic systems relies on the incident solar radiation, while the generation of wind turbines relies on the prevailing wind speed. This raises particular stability issues, as the demand is no longer able to follow a highly uncontrollable supply. Against this background, energy generation based on more dispatchable renewable generations such as hydroelectric or geothermal have been proposed. Also, the generation based on other fuels, such as nuclear power have also been proposed as a way forward.

An alternative approach to counter these issues is the introduction of Demand Side Management (DSM) incentivized policies and technologies. In more detail, DSM stands for the concept of having the supply side trying to match an uncontrollable demand. This can be achieved by incentive programs or variable tariffs.

For instance, a simple DSM policy is having prices that are higher during the periods of high energy consumption in order to incentivize the consumers to move their consumption during off peak hours. For such approaches to operate well, fast and rational control from the consumption side that respects the prices and the incentives provided by the grid operator needs to happen. As such, advanced economic control becomes an essential part of electricity consuming devices in the next generation electrical grid. Furthermore, advanced economic control can further increase the efficiency of such devices, as it can ensure an efficient operation. For instance, Heating, Ventilation and Air Condition (HVAC) systems can be ensured that they operate during occupation hours, in order to further increase the efficiency of the system, or they can be ensured that they do not overheat or over-cool the system.

In the context of the DSM schema, efficient HVAC control can preheat or precool a space before it is occupied if the electricity price is higher during the occupation period, and just let the temperature to drop/raise to the preference levels when the space is expected to be occupied.

To sum up, it is evident that in the new generation electrical grid efficient control of electrical devices becomes an integral part of for supporting an ever increasing energy demand and reducing on reliance on fossil fuels.

Due to these reasons, this thesis preoccupies itself with efficient control and in particular with efficient control of cooling, heating and ventilation. Buildings are considered the 60% of the electrical energy consumption in the EU (Marie Rousselot, Energy efficiency trends in buildings , Policy Brief 2018)[19], and HVAC is considered about 11% of the electrical energy consumption (REHVA Journal 01/2012)[18]. As such, increasing the energy efficiency of such devices, through efficient control, provides great opportunities for:

1. Increasing the energy efficiency.

2. Provide opportunities for Demand Side Management.

The latter is also supported by the fact that thermostatically controlled loads (such as heating control and ventilation) are considered a natural energy buffer, that allows energy to be stored in the form of heat until it used during occupation, when it is necessary. In this context, Demand Side Management i.e.: shifting the demand to non peak periods or periods that the grid operator needs the demand to be shifted to, can happen without the usage of energy storage technologies, such as batteries, flywheels etc, which can be particularly expensive.

In the next section, we provide more details about our work on HVAC control.

1.1 Heating, Ventilation and Air Conditioning (HVAC)

As discussed above, HVAC represents one of the biggest shares of today's residential and commercial buildings' energy consumption (EIA 2015)[3]. As such, working towards more efficient HVAC systems can lower the energy bills and carbon emissions. The use of optimization algorithms in this pursuit for higher efficiency can provide us with a highly cost effective HVAC operation with minimum disruption, especially when compared to more intrusive energy efficiency improvements, such as replacing the insulation of a building or installing new HVAC equipment. For this reason, optimizing the HVAC control process has been heralded as a key means for energy efficiency improvements in today's buildings (Dounis and Caraiscos 2009)[2], preparing the way for an energy sustainable future.

During the optimization of the HVAC control process, the most prevalent trade-off is the one between the thermal discomfort that the occupants experience and the energy cost required. Thus, this can be technically represented as a multi-objective optimization problem, with the two main and conflicting objectives being the occupants' thermal comfort preferences and the energy cost of the system. Besides these two objectives, one must also pay attention to the building safety (e.g., the minimum and maximum allowed temperatures) and any other regulatory constraints. In settings with a fixed occupancy or working schedule, such as in commercial buildings, the optimization of the HVAC control process can be achieved by planning for the minimum energy cost required to meet the thermal constraints that are set by the users. These constraints typically take the form of a temperature comfort band, defined by fixed setpoints of maximum and minimum temperatures, for particular time periods (e.g., the duration of a work day). However, in dynamic occupancy settings ,such as in residential buildings and office buildings where the occupants don't have a predetermined work schedule, this approach can be found inadequate since the optimization process can be hindered by occupancy estimates, that are inherently uncertain. More specifically, strictly maintaining the temperature of a room so that it lays within a narrow comfort band when there is only a small probability of occupancy can significantly increase the energy usage of the HVAC system. In this case, it is essential to alleviate the cost by accepting some probability of discomfort. Although, this approach creates the additional challenge of defining the amount of acceptable discomfort that the occupants could experience, according to their preferences.

A number of approaches for optimizing the HVAC control process have been proposed over time that deal with probabilistic occupancy estimates (Gao and Keshav 2013; Lu et al. 2010; Panagopoulos et al. 2015; Scott et al. 2011; Urieli and Stone 2013)[5][6][7][8][14]. The two main lines of research that derive from these works are the thresholding and the weighted sum approaches. The former relies on defining a probability threshold above which the occupancy is considered certain, thus thresholding the probabilistic occupancy estimates and optimizing the HVAC control process on the derived deterministic occupancy schedule. The latter utilizes the probabilistic occupancy schedules as it is and then relies on a weighting parameter to determine the balance between the HVAC energy cost and the expected thermal discomfort (Panagopoulos et al. 2015)[7]. Both of these approaches have pros and cons with respect to the criteria of effectiveness, efficiency, applicability and usability. Nevertheless, these advantages and disadvantages have not been thoroughly understood and evaluated. Importantly, the critical question of how to meet the user preferences in balancing heating cost and thermal discomfort is usually ignored in the literature. A notable exception, Panagopoulos et al. 2015 is our point of departure and provides preliminary results in this direction.

In the context of small commercial buildings, we identify the following necessary requirements for a suitable balancing technique:

1. **Applicability**: It is crucial for the developed method to be able to be applied in today's smart thermostats and operate within the computation resources requirements (for instance cloud services are not usually available in domestic settings and their usage would render such an approach cost ineffective). Also, it should be able to operate while meeting the real time operation requirements of such systems. Furthermore, it should rely to minimum extent to instrumentation (as extensive sensor networks are not usually available in domestic settings) and additional information.

2. **Effectiveness**: Our proposed approach should be able to balance thermal discomfort and HVAC cost in order to meet the user preferences.

3. **Efficiency**: Our proposed approach should be able to balance thermal HVAC cost and thermal discomfort with maximum efficiency. In a bi-objective optimization problem as this one, efficiency means that the approach should capture solutions on the Pareto frontier. Dominated solutions should not be returned by the proposed approach (i.e., solutions that can achieve a particular cost with lower discomfort or a certain discomfort with lower cost should not be returned as final solutions). Returning dominated solutions means that you can achieve the same discomfort or the same cost with lower cost or lower discomfort respectively.

4. **Usability**: It is critical for the proposed approach to be able to respect the user preferences, as discussed in the effectiveness criterion. However, this should also be possible through a user friendly procedure. For instance, populating complex and hard to comprehend is not always the most user friendly solution.

1.2 Thesis Contribution

In this thesis we provide a qualitative and quantitative comparison of the 2 approaches, mentioned in the previous section, evaluating their efficiency, effectiveness, applicability and usability. We also propose a new approach based on variable bounding and show that it is capable of capturing optimal solutions with minimum user input. Our study shows that the weighted sum formulation is a clear winner in terms of the range of user preferences that is able to capture. All approaches, however, capture optimal solutions with minimum user input. As such, the preference of one approach over the others depends on the specifications of the application with respect to user input. In this work we also show that the setpoint temperature set by the user should consider the origin of the discomfort metric in stochastic occupancy settings, i.e., the region where the occupant feels absolute thermal comfort, and not a parameter to balance discomfort and cost as this balancing must occur while respecting the expected occupancy probabilities.

The work in this thesis was a part of selecting the most suitable algorithm for balancing HVAC operation cost and expected occupant thermal discomfort in real world trials. These real world trials were conducted as a part of the XBOS-DR: Customer- Controlled, Price Mediated, Automated Demand Response for Commercial Buildings project, that had as its main objective the development of an operating system for buildings that could support various intelligent applications. This project was a collaborative research with partners from UC Berkeley, Technical University of Crete, Siemens and Quest and was funded by the California Energy Commission. Preliminary results of this trial are also included and discussed int this work. Moreover, the work in this thesis led to a publication that appeared in the Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings, 2018. [1]

1.3 Thesis Outline

The rest of the thesis is structured as follows: In Chapter 2 we provide the needed background material. Following, in Chapter 3 we investigate two approaches for dealing with expected discomfort, as well as introducing a new one. In Chapter 4 we discuss the details of our experimental setups. In Chapter 5 the evaluation results are discussed. Chapter 6 discusses some aspects of the aforementioned real world trials and finally Chapter 7 provides conclusions and outlines future work directions.

Chapter 2

Related Work and Background Material

2.1 Model Predictive Control

Model Predictive Control (MPC) is a wide family of approaches for the control of dynamic systems. These approaches are based on prediction over the state of a system over a finite time horizon, computing possible future inputs at each step by minimizing a cost function and controlling the system accordingly (Camacho & Bordons, 2004)[17]. MPC is used on the state-of-the-art HVAC control process optimization approaches, formulating it into the following steps (Camacho and Alba 2013):

1. After a predetermined time interval, plan a HVAC control schedule over a finite horizon into the future using models of the system dynamics.

2. Execute the first action of the planned schedule.

3. Repeat the procedure by shifting the planning horizon into the future.

In order to create control schedule models for thermal dynamics, thermal discomfort, local weather, occupancy prediction (in the case of a building with dynamic occupancy schedule) and energy cost must be created. In some cases a renewable generation model and an energy consumption model of the rest of the building loads can be considered as well.

In HVAC control, the parameters of the learned models can be updated after each execution. This variant, called Adaptive Model Predictive Control, can introduce stability issues in some settings, but in the case of HVAC control the thermal dynamics are slow enough, so this approach can be considered without considerable respective concerns (Siroky et al. 2011)[9].

2.2 Thermal Modeling

As discussed above, modelling is considered an integral part of Model Predictive Control, as a model is utilized to plan an action schedule. In the context of HVAC Model Predictive Control it is essential to have a thermal model of the building. In more detail, a thermal model considers a function that aims to predict the future thermal state of the building, given the current building thermal state and/or additional information, along with the HVAC control action, such as heating or cooling, or the setpoint temperature. As in all modeling approaches, thermal modeling can also be divided into three broad categories (Prívara et al. 2013)[10]:

- White-box
- Gray-box
- Black-box

White-box modelling refers to the concept of utilizing all available physical laws (or as much as possible), in order to make an extensive description of the thermal dynamics of a building. For instance, a white box model would take into account the thermal resistance of each one of the materials that comprise the walls, each one of the materials that are used for the windows and make a complex network of the thermal dynamics of the building. Usually, these thermal dynamics are modelled using resistance capacitors networks, that utilize an analogy with electric networks. However, the main drawback of such approaches is that they require extensive knowledge about the building and/or blueprints of the building, which are not readily available in most cases, especially in older buildings.

On the other hand, black-box approaches learn the thermal response of a building utilizing only recorded data that can be collected from the building, and as such overcome the limitations of white-box approaches. Black-box approaches include thermal modelling approaches based on neural networks, Support Vector Machines, polynomial curve fits and other machine or statistical learning techniques (e.g., Ruano et al. 2006, Huang et al. 2013)[11][12]. Nevertheless, the main limitation of such approaches is that the black-box utilized doesn't have a physical meaning, and as such it might predict non physical values (i.e.: a neural network might be predicting negative relative humidity values). This is especially the case if the black-box approach is undertrained, because not enough data is available, or the data is not covering the whole region of operation evenly.

Against these two approaches stand gray-box approaches, which are hybrid approaches aiming to overcome both limitation by combining both previous approaches. Gray-box approaches rely on simplified physical equations, where the equivalent thermal parameters are being trained using data. As such they retain some physical meaning, while not requiring extensive knowledge that is not available. Due to these reasons, in this work a gray-box approach is utilized and in particular we utilize a simple thermal model described below.

To create a thermal model of a building, an HVAC control action, a set of variables (e.g., outside temperature, incident solar radiation), a certain time horizon, and the current thermal state of the building serve as input to a function that links them to an approximate future thermal state after the given time horizon has passed. More formally the thermal model can be expressed as:

$$x_{t+1} = TM(x_t, a, t, i)$$

where **x** is the thermal state vector. **a** is the HVAC control action vector, **t** is the time horizon and **i** is the additional variables vector.

2.3 Predicting Occupancy

Predicting the occupancy schedule in buildings with dynamic occupancy is a crucial task that supports a number of intelligent applications in the context of smart buildings, such as optimizing the charging of electric vehicles, ensuring safety and security requirements, fault diagnostics etc. In this context, a number of approaches have been proposed for predicting occupancy schedules; these approaches rely on different input signals such as passive infrared PIR sensors, GPS signals from smartphones, WIFI usage etc.

In more detail, occupancy prediction approaches can be classified in two broad categories, historical data-based and context aware. Historical data-based approaches rely only on the past occupancy schedule to make predictions. On the other hand, context aware approaches utilize also information about the context of the occupants, such as their position and their proximity to a target zone.

The former approaches require less instrumentation and hence they are appropriate for the majority of today's buildings and require minimum retrofitting. In contrast, the latter approaches require extensive instrumentation that is not typically available in today's buildings. Nevertheless, despite the lack of information, the historical data-based approaches have demonstrated high accuracy in predicting the occupancy schedule (around 90% as demonstrated in (J. Scott et al. 2011))[8].

In this work, we focus on such (i.e., historical data-based) approaches, and in particular in the approach proposed in (J. Scott et al. 2011) [8], that produces state of the art predictions in comparison to other such approaches.

2.4 Thermal Discomfort

As discussed, modelling is considered an integral part of MPC, and along with thermal modeling it is crucial to have a model of also the thermal discomfort of the occupants. In more detail, thermal discomfort aims to capture the deviation of the thermal state of the building, from a thermal state that is considered the optimal with respect to the user preferences.

Thermal discomfort is generally a non-tractable concept that is affected by various factors, ranging from the temperature of the building to relative humidity, the activity of the user, the wind speed, the mood of the user, clothing and more. As such, various thermal discomfort metrics have been proposed that aim to quantify thermal discomfort utilizing such factors. One the most prominent thermal discomfort modelling approaches considers the ASHRAE thermal sensation scale, that standardizes the thermal discomfort into seven (7) different values, taking into account the aforementioned factors.

The adaptive ASHRAE model has also been proposed for European buildings. The adaptive ASHRAE model relates indoor design temperatures or acceptable temperature ranges to outdoor meteorological or climatological parameters (ASHRAE Standard 55, 2013). [13]

That said, taking into account all these factors when modelling thermal discomfort is extremely challenging, as this would require extensive sensors and instrumentation to detect relative humidity, clothing, activity or even mood. As such, simplified versions of this thermal discomfort model are being utilized in practice.

In this work, for simplicity, we are utilizing a thermal discomfort model that penalizes discomfort when it deviates from a user provided comfort zone. In more detail, if the temperature is below or above user provided temperatures, we quantify discomfort as the amount of time that actually the temperatures were outside of this region. For instance, in Figure 2.1 the quantified discomfort would be indicated by the area that is highlighted by the gray colored region, which corresponds to the deviation of the temperature from the user provided comfort-band.

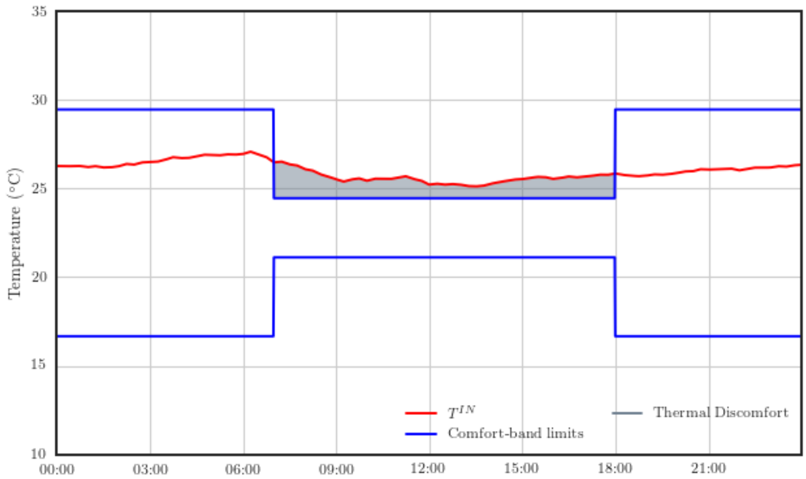

Figure 2.1: Thermal Discomfort when the temperature is outside the comfort-band limits.

Arguably, any other discomfort thermal metric can be used for our work here, without loss of generality, as the comparison of all the approaches that we have utilize exactly the same thermal discomfort metric.

A quantitative metric of the discomfort experienced by the occupants relies on a user-provided comfort band in the form of thermostat set-points. These set-points create a comfort band, that must be guaranteed while the building is occupied, in order for the occupants to feel no thermal discomfort.

A simple discomfort metric that highly penalizes thermal discomfort is:

$$Disc = \begin{cases} (T^{in} - T^{up})^2 & \text{if } T^{in} > T^{up} \\ (T^{in} - T^{down})^2 & \text{if } T^{in} < T^{down} \\ 0 & \text{otherwise} \end{cases} \tag{2.1}$$

where $Disc$ is the instantaneous thermal discomfort that occupants experience, T^{in} is the indoor temperature and T^{up}, T^{down} are the upper and lower limits of the comfort band, respectively. This is the metric that we adopt in our work here.

2.5 Expected Thermal Discomfort

Thermal discomfort can only be experienced when the building is occupied, and since occupancy predictions can be incorporated while planing the optimal HVAC control schedule, the expected thermal discomfort that an occupant experiences is calculated as:

$$E[Disc] = \begin{cases} O(T^{in} - T^{up})^2 & \text{if } T^{in} > T^{up} \\ O(T^{in} - T^{down})^2 & \text{if } T^{in} < T^{down} \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

where $E[Disc]$ is the expected instantaneous thermal discomfort that the occupants are expected to experience, O stands for the probabilistic occupancy estimates, T^{in} is the indoor temperature and T^{up}, T^{down} are the upper and lower limits of the comfort band, respectively. This is the metric we adopt in our work here.

Chapter 3

Dealing with Expected Thermal Discomfort

As discussed in Chapter 2, MPC deals with optimizing control processes, especially infinite horizon ones, when a model is available. Such an optimization task can be formulated as a receding-horizon planning, which repeatedly plans an optimal action schedule ahead, according to some objective. That said, in the context of HVAC, the optimization considers a multi-objective one. In more detail, the objectives that one has to optimize in the context of HVAC consider cost, thermal discomfort and also other secondary thermal requirements, including safety limits such as not exceeding the temperature above particular values so that the building is not damaged.

Two of these objectives are core, and, in this context, we can regard HVAC as a bi-objective optimization (including cost and thermal discomfort) and also certain constraints that are being put on top of that. In general, in multi-objective optimization the aim is to find solutions that fall into the Pareto frontier, the set of all Pareto optimal allocations (Osborne and Rubinstein, 1994, p.122)[21]. In more detail, Pareto optimality considers the concept where none of the objective functions can be improved in value without degrading the value of the other objective functions. Moreover, the solution that fall into the Pareto frontier are the ones that can achieve an optimal allocation of resources, and the solutions that do not fall into the Pareto frontier are called Pareto dominated solutions. As such, in the context of HVAC we also want to find a solution that balances cost and discomfort falling on the Pareto frontier. In this context, bi-objective optimization can be formulated either as a single objective optimization, which constraints on the second objective, or combine both objectives through a unifying function, such as the scalarized weighted sum function. The first approach, optimizing one objective while keeping the other as a constraint, is computationally very expensive. As such, traditionally such applications focus on unifying approaches.

In this work we are considering and evaluating (as discussed in Chapter 1) two well known unifying approaches and one new that we propose. In particular, the basic challenge of HVAC control process optimization approaches considers dealing with expected discomfort (i.e. step I in the MPC procedure above) and, in particular, on how the HVAC control schedule is calculated given this uncertainty. In this context, one can distinguish two main lines of thought: the thresholding approach, and the weighted sum approach, which we detail in the following paragraphs. Then we detail a new variable bounding approach that we propose.

3.1 Thresholding

Thresholding the probability of occupancy can be achieved by arbitrarily selecting a threshold, which is then used for deriving a deterministic schedule of occupancy. A common threshold of choice has a value of 0.5, where all the probabilities of occupancy that have greater values than this threshold are pushed to 1. Likewise, all occupancy probabilities below this threshold are pushed to 0. A typical probabilistic occupancy estimate vector can be seen below:

$$\begin{bmatrix} 0.2 & 0.4 & 0.5 & 0.6 & 0.2 & 0.5 & 0.5 & 0.4 \end{bmatrix}$$

were each number corresponds to a particular time interval (e.g., 15 minutes) within a day. By setting the threshold to 0.5, the occupancy schedule is turned from stochastic into a deterministic vector where 1 corresponds to certainty of occupancy, while 0 corresponds to certainty of vacancy:

$$\begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}$$

Subsequently, the optimization algorithm takes into account the user-provided comfort bounds only when occupancy is expected (i.e., time intervals where the deterministic occupancy is equal to 1), and respects the safety requirements or does no action otherwise. In doing so, these methods do not cool or heat the room when there is a low likelihood of occupancy.

However, the algorithm will have the same behaviour for time intervals with probabilistic occupancy estimates that are very close to 0 as well as for time intervals where the probabilistic occupancy estimates are just below the threshold (e.g., 0.49 when a threshold with a value 0.5 is selected). This is not optimal in many situations, since the optimization algorithm should preheat/precool the thermal zone in order for it to be closer to being comfortable when the probability of occupancy is higher. Despite these weaknesses, this approach is intuitive and simple enough to allow for straightforward human-computer interaction. In particular, this method allows the user to visualize the anticipated deterministic occupancy schedule, thus letting him/her to make informed decisions about overriding the intelligence. Although the user may adjust the threshold, the intended user involvement is to only populate the comfort bands (e.g., the heating and cooling setpoints) for each day, which is a straightforward task.

3.2 Weighted Sum Balancing

Weighted sum balancing is another approach to deal with expected thermal discomfort. The Weighted Sum Balancing approach utilizes probabilistic occupancy schedules and also a weighting parameter in order to balance the expected thermal discomfort and the heating/cooling cost. In this context, the comfort constraint is integrated in the objective function. In more detail, most of these approaches minimize the weighted sum, J, of cost and discomfort as expressed in the following form:

$$J = (1 - \lambda)Cost + \lambda E[Disc]$$

where λ is a weighting parameter which ranges from 0 to 1. Note here, that 0 and 1 are excluded since the other objective is omitted. If λ lies closer to 1 then discomfort is valued more compared to cost. In contrast, when λ falls closer to 0 then cost is valued more compared to discomfort.

Such scalarized approaches, that unify cost and discomfort in a single objective, deal with the probabilistic occupancy estimates in a mathematically concrete manner. In this context, they provide Pareto optimal solutions. Pareto optimal solutions have an optimal balance of cost and discomfort. This means that the same cost of a particular solution cannot be achieved with less discomfort and vice versa. Notably, identifying the value of λ that perfectly captures the user preferences is a hard task. This is the case, because cost and discomfort are measured in different units. Moreover, the relationship between cost and discomfort is hard to be understood and interpreted by the occupants. That said, given that λ is a single parameter, an iterative approach of populating it can be used in order to meet the user preferences as illustrated in Panagopoulos et al. (2015)[7].

Now, hybrid approaches also exist that utilize both a thresholding and a weighted sum balancing formulation, such as SPOT+ by Gao and Keshav (2013)[5]. However, these fall short in all categories of efficiency, usability and applicability compared to the weighted sum balancing or the thresholding approach (Panagopoulos et al. 2015)[7]. As such, these approaches are not followed in this work.

3.3 Variable Bounding

Here, we propose a new approach to deal with expected thermal discomfort. This approach relies on variable bounding. In more detail, in our approach the comfort-bands are adjusted in a dynamic manner and in accordance to the probabilities of occupancy. Hence, the comfort-band bounds are closer to the user provided values when the occupancy probabilities are closer to 1, while the bounds become wider as the probabilities get closer to 0. As such, as the probability of occupancy becomes smaller, a greater deviation from the original comfort-band is allowed. That said, it is not trivial to identify how wide should the bounds become in accordance to the occupancy probability. One can suggest that the bounds should become infinitely wide with occupancy probabilities close to 0. However, a more narrow practical limit to the comfort-band boundaries could be held useful in order to avoid unnecessary uncomfortable indoor thermal conditions and facilitate a linear expansion function. More formally, the following formula can be used to relate the expansion of the bands to the occupancy probabilities:

$$T^{up'} = OT^{up} + (1 - O)T^{maxup}$$

$$T^{down'} = OT^{down} + (1 - O)T^{mindown}$$

where $T^{up'}$ and $T^{down'}$ stand for the upper and lower expanded band limit respectively while T^{maxup} and $T^{mindown}$ stand for the upper and lower absolute limits, respectively. These upper and lower limits can also consider the safety thermal requirements of the building.

Chapter 4

Experimental Setup

For our case study we consider a small building with 4 distinct thermal zones in Berkeley, California USA. We choose one zone for our evaluation, namely $T2$. The zone is instrumented using wireless occupancy and temperature sensors which collect data and send them to a local server running the building operating system described in Fierro and Culler (2015)[4]. Our approach is explained further in this chapter.

4.1 Case Study and Discomfort Evaluation

For our evaluation we use data for one day during the winter and evaluate our approach for this particular day through an iterative procedure. In this procedure we ensure that the thermal state at the beginning and end of the day is the same. In this way, assuming that the same day repeats itself, we are able to provide long term evaluation results in feasible time. This enables us to evaluate all of the considered approaches with a wide population of the corresponding balancing parameter (i.e. the threshold, the λ parameter, and the T^{maxup}, $T^{mindown}$ values, respectively). In more detail, all approaches were evaluated using the following ranges, for the corresponding parameters:

1. Thresholding: A threshold within the 0-1 range with a step of 0.1.
2. Weighted Sum Balancing: within the 0-1 range with a step of 0.02.
3. Variable Bounding: T^{maxup} and $T^{mindown}$ from the comfort band limits to 4 times the safety requirements with a variable step of 0.1 to 1.

Starting at the beginning of the chosen day, the optimal action is acquired from the MPC planning algorithm and it is simulated through the thermal model for a 15 minute time interval, giving us the new thermal state of the zone. To this new thermal state, a correction bias is added, sampled from a Gaussian Distribution with standard deviation and mean equal to the standard deviation and mean of the error of the thermal model for the 15 minute time interval. Then time is shifted by 15 minutes and the same process repeats for the whole day.

4.2 Prediction Models

Occupancy is predicted using a state-of-the-art algorithm proposed by (Scott et al. 2011) that reportedly achieves $\sim 80\%$ accuracy.

This is a similarity based approach, using 60 days of historical data to find the 10 days that are the most similar to the day we are examining. This similarity is based on the already known occupancy of the day, adding 4 hours from the previous day to avoid not having enough data at the start of the day. Then the occupancy schedule is predicted as an average of the similar historical days.

In this work we also evaluate this algorithm for our zones. As can be seen in Figure 4.1 the occupancy prediction algorithm utilized achieves prediction accuracy in the range of 70-90%, which is consistent with the previous evaluations.

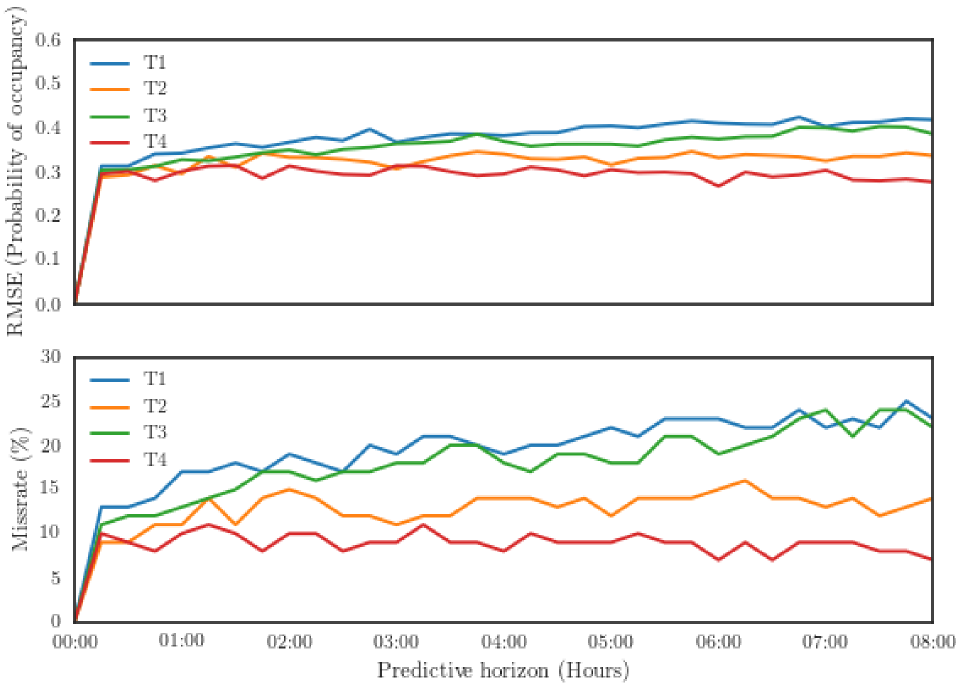

Figure 4.1: Occupancy prediction approach evaluation for all zones.

Thermal response is predicted using a simple linear formulation and, in particular:

$$T^{in}_{t+\Delta} = T^{in}_t + (c_1 T^{in}_t h + c_2 T^{in}_t c + c_3(T^{out}_t - T^{in}_t))\Delta$$

where, $T^{in}_{t+\Delta}$ is the estimated indoor temperature after Δ amount of time while T^{in}_t and T^{out}_t are the current indoor and outdoor temperatures respectively. The parameters c_1, c_2 and c_3 are the regression coefficients to be estimated. The thermal model is estimated through online least-squares fitting regression. The predictive accuracy of the model is reported in Figure 2. Outdoor temperature predictions were acquired through online meteorological providers and in particular wunderground (https://www.wunderground.com).

In this work we also evaluate the thermal model prediction accuracy, to show that the root mean square error (RMSE) for an 8 hour ahead prediction is about 2.5F which is consistent with the state-of-the-art. (As can be seen in Figure 4.2)

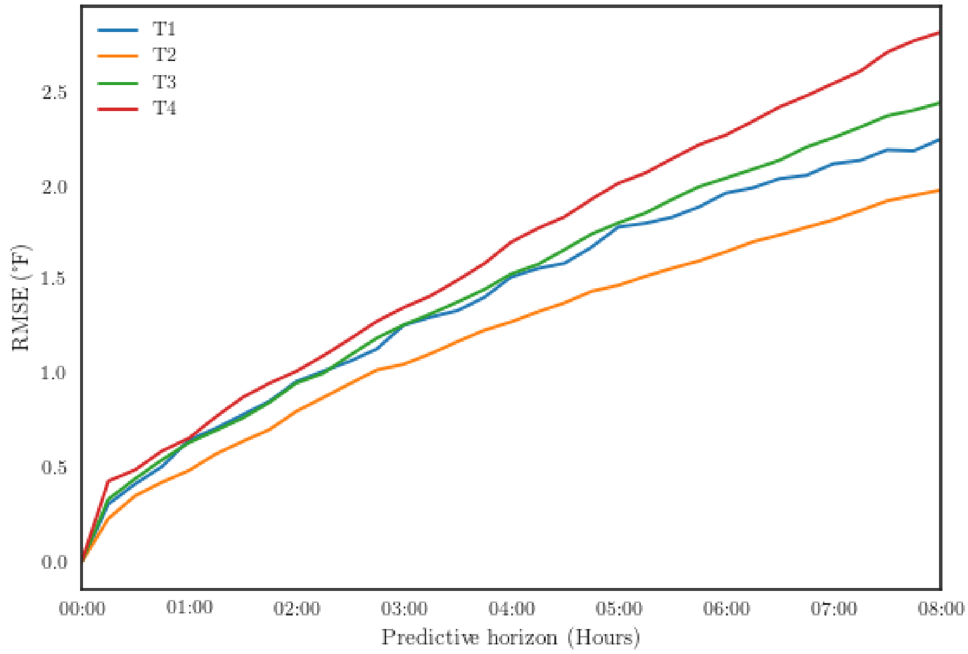

Figure 4.2: Thermal model prediction evaluation for all zones.

4.3 Model Predictive Control

The control schedule for the HVAC is handled with an MPC algorithm by creating a Directed Acyclic Graph (DAG), where the starting node contains all the info about the present state of the thermal zone. Each node points to at most 3 other nodes, created by simulating the heating, cooling or idle actions through the aforementioned models.

Adaptive MPC could be used on a real life scenario (such as the real-world trials discussed in Chapter 6), since the thermal dynamics of buildings tend to be slow. Though, because the data created through the simulation can incorporate cumulative errors if used to retrain the models, for the purpose of these experiments we did not use the Adaptive MPC approach.

While creating the DAG to plan the HVAC control schedule, after simulating the same actions, but with different ordering (i.e., Heating followed by Idle and

Idle followed by Heating actions), the nodes created can predict almost identical thermal states for a given time interval. Rounding up the zone temperatures inside the graph generation algorithm can speed up the execution time significantly, by creating one single node for nodes with a very similar thermal state at the same time interval. Trading off precision for computational efficiency in this setting is very important, because the control schedule must be created before the next time interval in a real life application. This could also be supported by an Approximate Dynamic Programming approach later on.

Some sample graphs created during the experiments can be seen in figures 4.3, 4.4, 4.5:

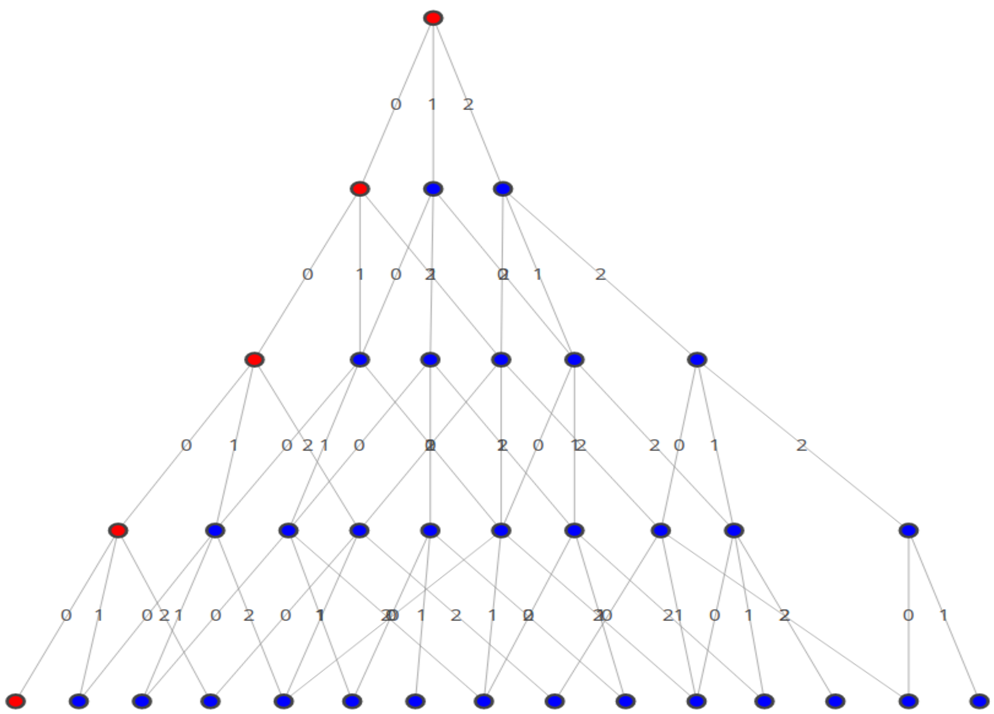

Figure 4.3: MPC graph created with a 1 hour planning horizon.

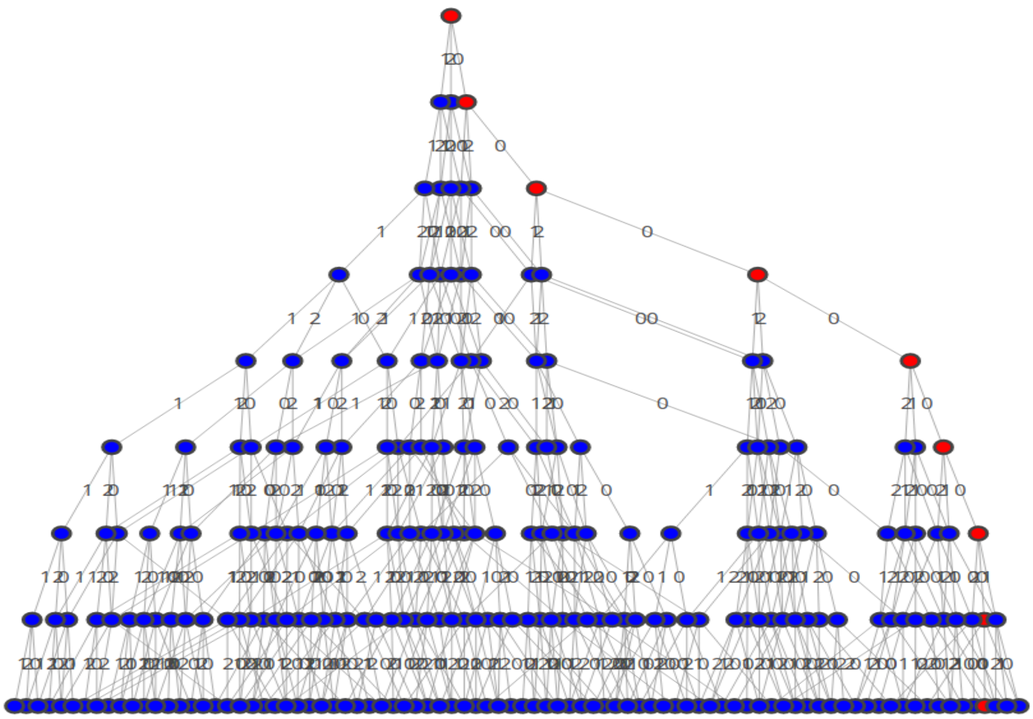

Figure 4.4: MPC graph created with a 2 hour planning horizon.

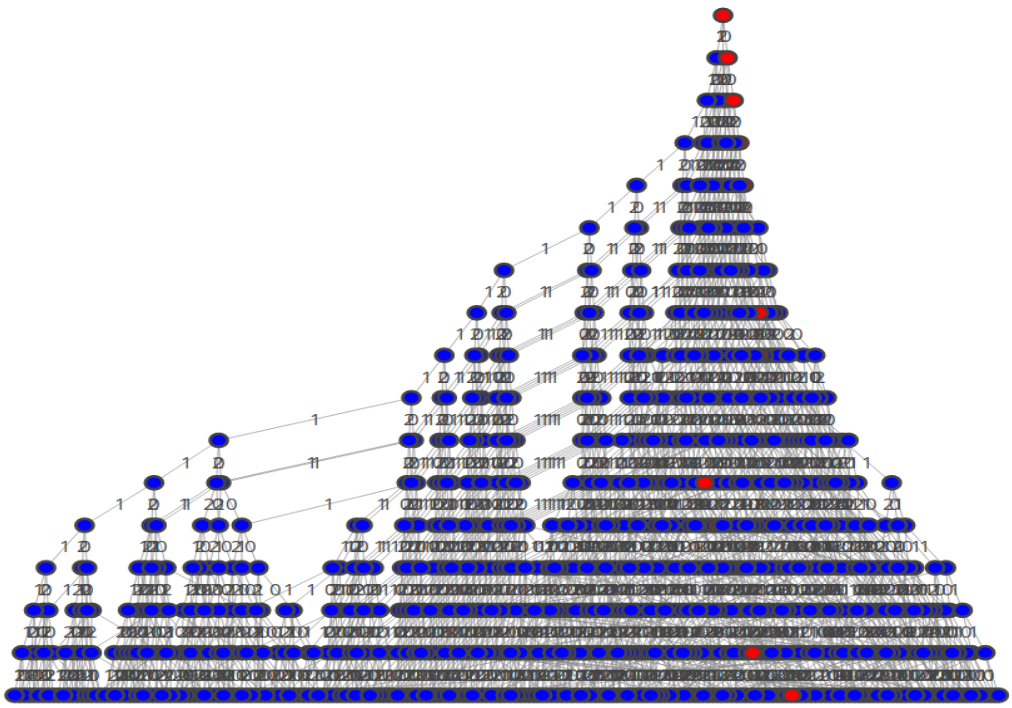

Figure 4.5: MPC graph created with a 4 hour planning horizon.

The nodes colored with red indicate the solution found by a shortest path dynamic programming algorithm. The numbers 0, 1 and 2 indicate the Idle, Cooling and Heating actions respectively. The first action of the optimal path is realized and the MPC planning algorithm is run again for the next time interval.

For the experiments, a planning horizon of 4 hours was used and the temperatures of the thermal zone of each node were rounded down to the first decimal. As you can see, the size of the tree rises exponentially with the planning horizon. The buildings safety constraints are implemented by not allowing node creation for temperatures that exceed these limits. In the case the algorithm runs for a starting thermal state that exceeds the safety constraints, the action that will bring the thermal state back to a safe state quicker is chosen.

Chapter 5

Evaluation Results

In this chapter we provide evaluation results and a thorough discussion. In more detail, the following paragraphs provide the evaluation results for each one of the balancing approaches, i.e. the Thresholding approach, the Weighted Sum Balancing approach and the Variable Bounding approach. We discuss our evaluation results with respect to the objective requirements as introduced in Section 1.1, in particular applicability, effectiveness, efficiency and usability.

5.1 Thresholding

In this section we discuss our evaluation results against the general requirements stated in Chapter 1 (i.e., applicability, effectiveness, efficiency, usability) for the Thresholding approach. Firstly, the Thresholding approach is an applicable approach, as it doesn't require extensive instrumentation and retrofitting and the evaluation results demonstrate that the approach can operate in real time as it is able to calculate a solution in milliseconds in a typical personal computer. Furthermore, Figure 5.1, 5.2, 5.3 and 5.4 illustrate the evaluation results of the approach with respect to different threshold parameters. In more detail, figure 5.1 provides a 3D illustration of the cost and discomfort against the threshold parameter, while figures 5.2, 5.3 and 5.4 provide illustrations of cost against discomfort, discomfort against the threshold parameter and cost against the threshold parameter for all different selected parameters respectively.

As can be seen the solutions captured by the Thresholding approach seem to fall on the Pareto frontier, while not any dominant solutions were captured. Hence, the approach meets the effectiveness and efficiency requirements as discussed in Section 1.1. Nevertheless, the Thresholding approach seems to capture only a narrow region of the Pareto optimal solutions as all the captured solutions range between \$1.27 cost and \$1.32 cost. Hence, not a wide range of user preferences can be met. Nonetheless, the approach relies only on a single parameter (the threshold parameter), hence in this context it is a user friendly approach. In addition, in the case where the users want to just have solutions that fall very close to minimum cost, the threshold parameter can be set fixed requiring as such not any parameter population from the users.

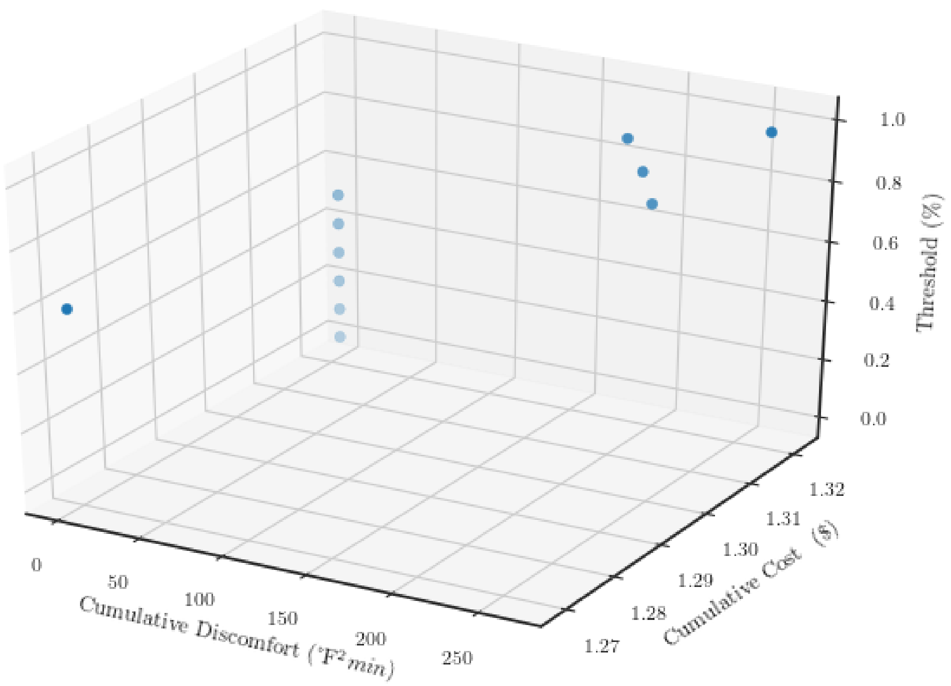

Figure 5.1: Aggregate Chart for the Thresholding approach.

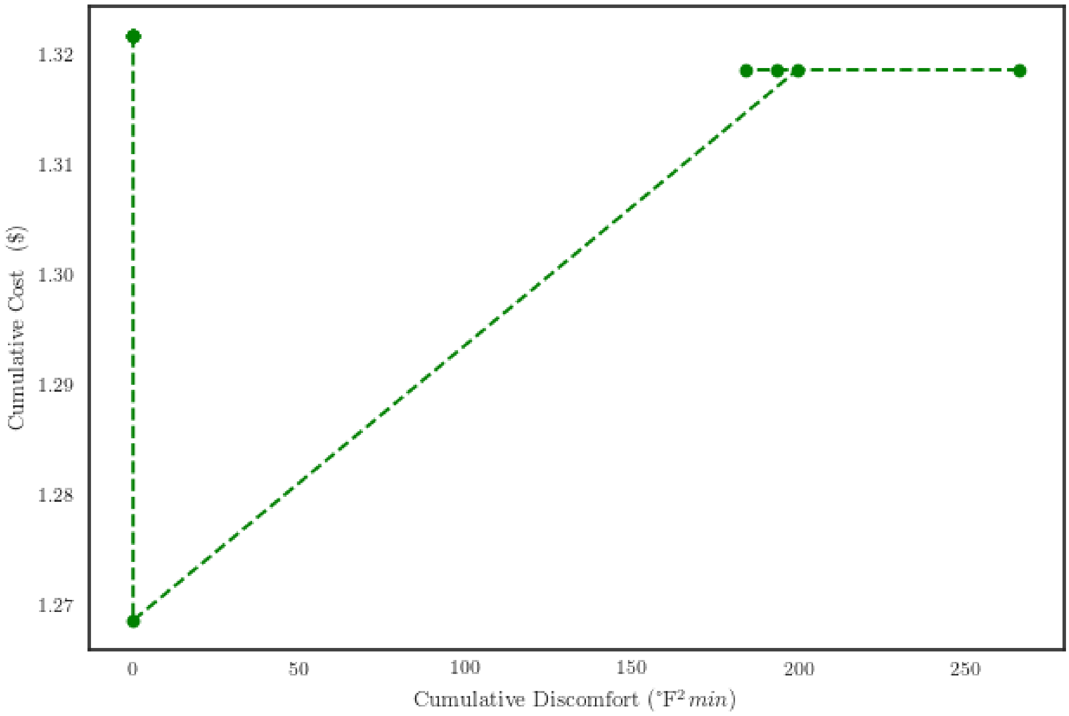

Figure 5.2: Balancing Cost and Discomfort for the Thresholding approach.

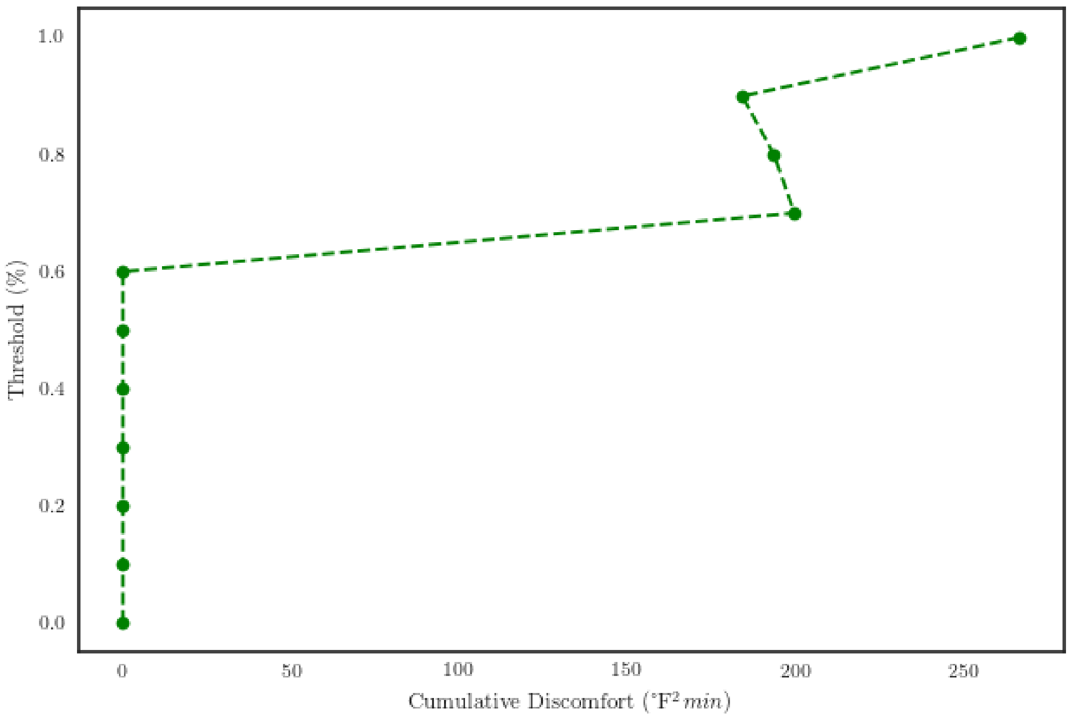

Figure 5.3: Balancing Discomfort for the Thresholding approach.

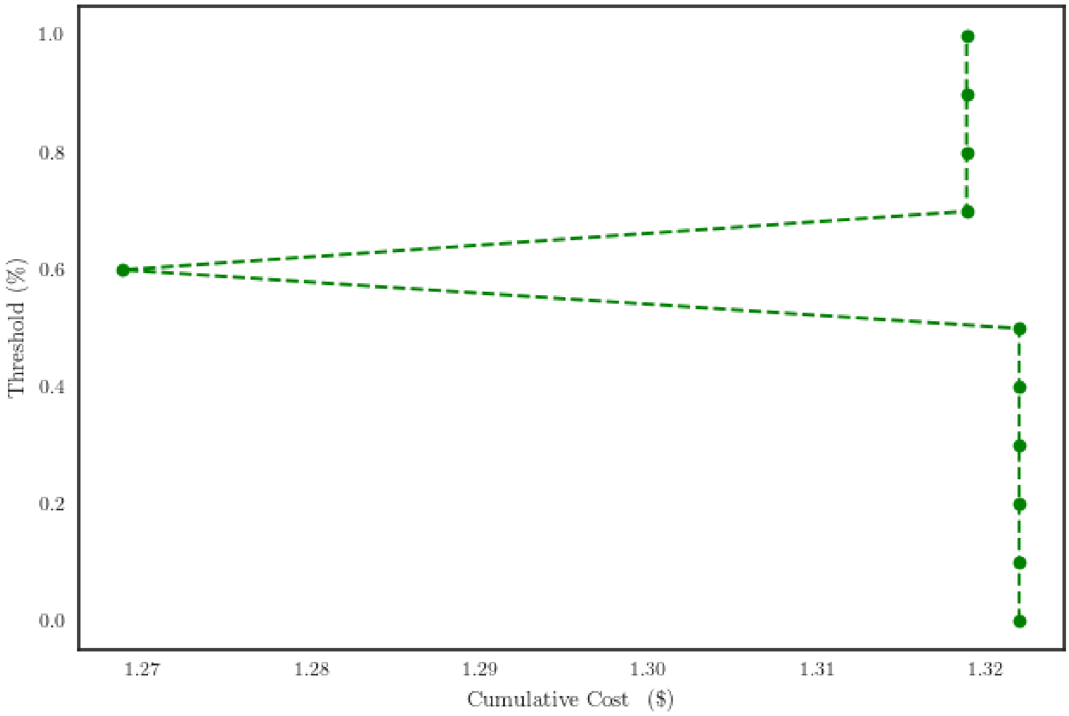

Figure 5.4: Balancing Cost for the Thresholding approach.

5.2 Variable Bounding

In this section we discuss our evaluation results against the general requirements stated in Chapter 1 (i.e., applicability, effectiveness, efficiency, usability) for the Variable Bounding approach.

First and foremost, with respect to applicability, also the Variable Bounding approach requires minimum computational time in a typical personal computer and hence we believe that it is totally applicable with even lower resource availability. It also requires minimal instrumentation and connectivity. In this context, this approach meets the applicability requirements as discussed in Section 1.1.

Now, figures 5.5, 5.6, 5.7 and 5.8 illustrate the evaluation results of this approach for different bounding parameters. In more detail, Figure 5.5 provides a 3D illustration of the cost and discomfort against the variable bound parameter, while figures 5.6, 5.7 and 5.8 provide illustrations of cost against discomfort, discomfort against the variable bound parameter and cost against the variable bound parameter for all different selected parameters respectively.

As can be seen, the approach effectively balances thermal discomfort and operating cost and hence meets the effectiveness requirement. With respect to efficiency, also this approach seems to capture solutions in the Pareto frontier and we were not able to identify considerably dominated solutions. Nevertheless, with respect to usability this approach also only captures a narrow range of Pareto optimal solutions and hence it is potentially not able to respect a wide range of user preferences with respect to balancing operational cost and thermal discomfort. For instance, if a user wants the operational costs to fall lower than the 1.29$ price no solution would be available. Nonetheless, if this approach is to be used in settings where the users wants to minimize their discomfort with minimal cost, this approach is also suitable and user friendly as the user is not required to populate any parameters, since any variable bound parameter will give solutions that are close to the same region.

In this context we advise the user to use a variable bounding parameter with a value between 0 and 0.5, since these values provide the minimum discomfort (as can be seen in Figure 5.7). Of course, if the user wants to minimize the operational cost then the heating system can be completely switched off. As such, this solution, although illustrated in the figure, does not add to the range of the solutions that the approach is able to capture (this solution is trivial).

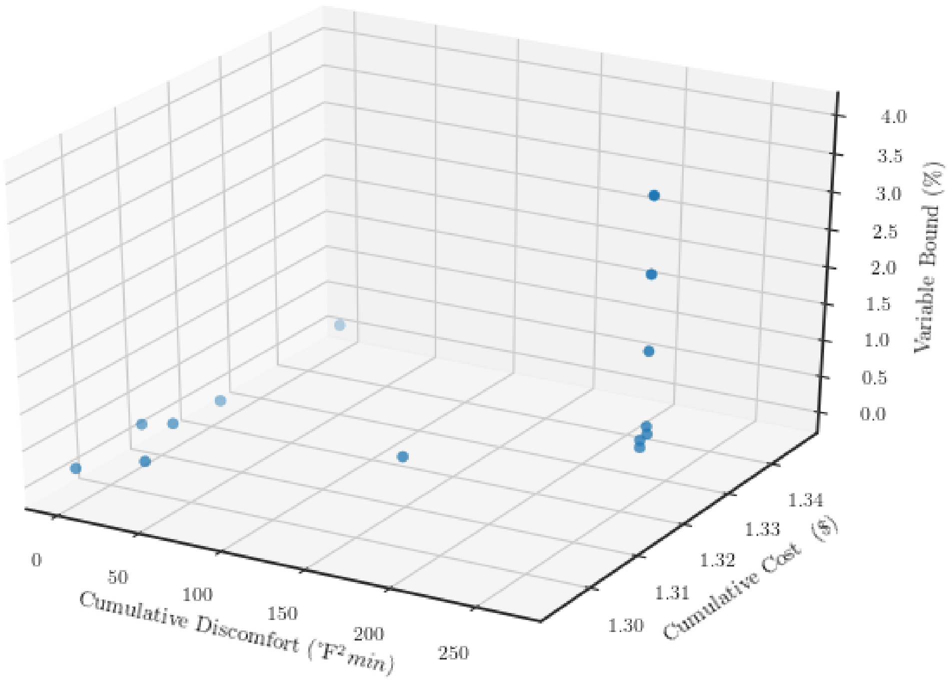

Figure 5.5: Aggregate Chart for the Variable Bounding approach.

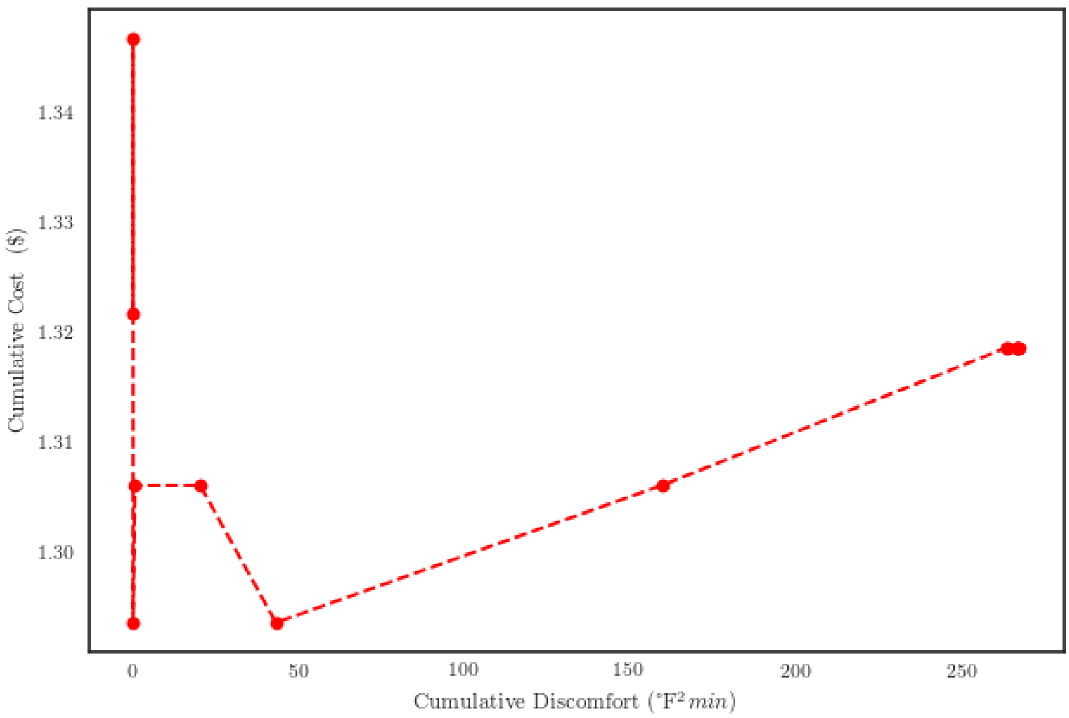

Figure 5.6: Balancing Cost and Discomfort for the Variable Bounding approach.

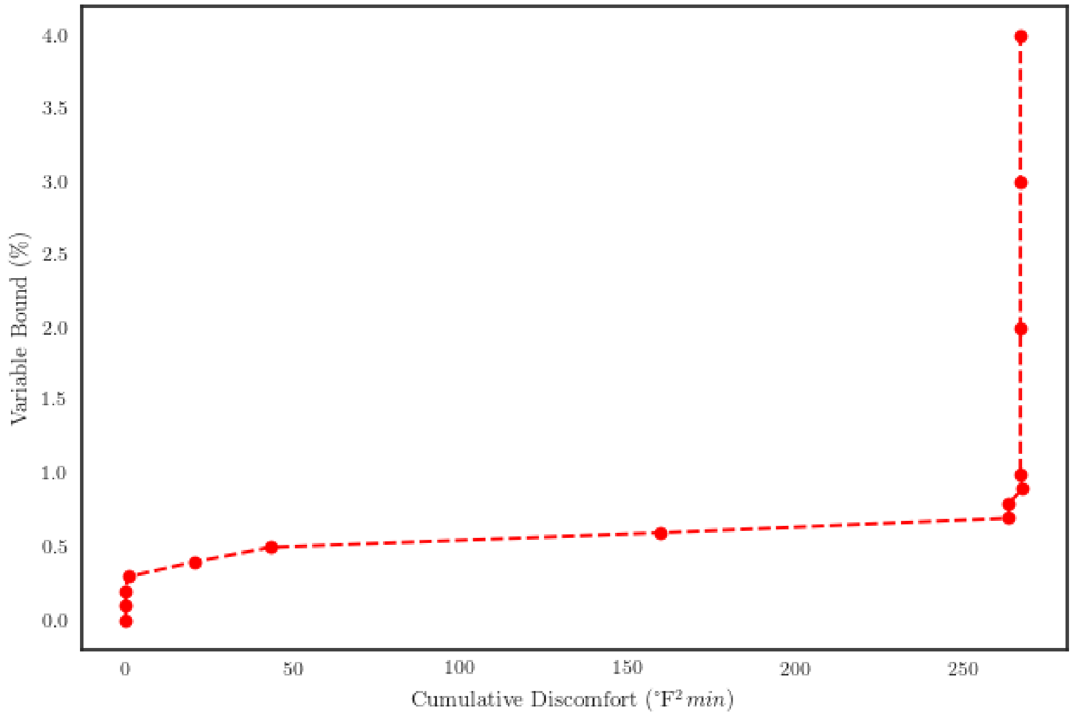

Figure 5.7: Balancing Discomfort for the Variable Bounding approach.

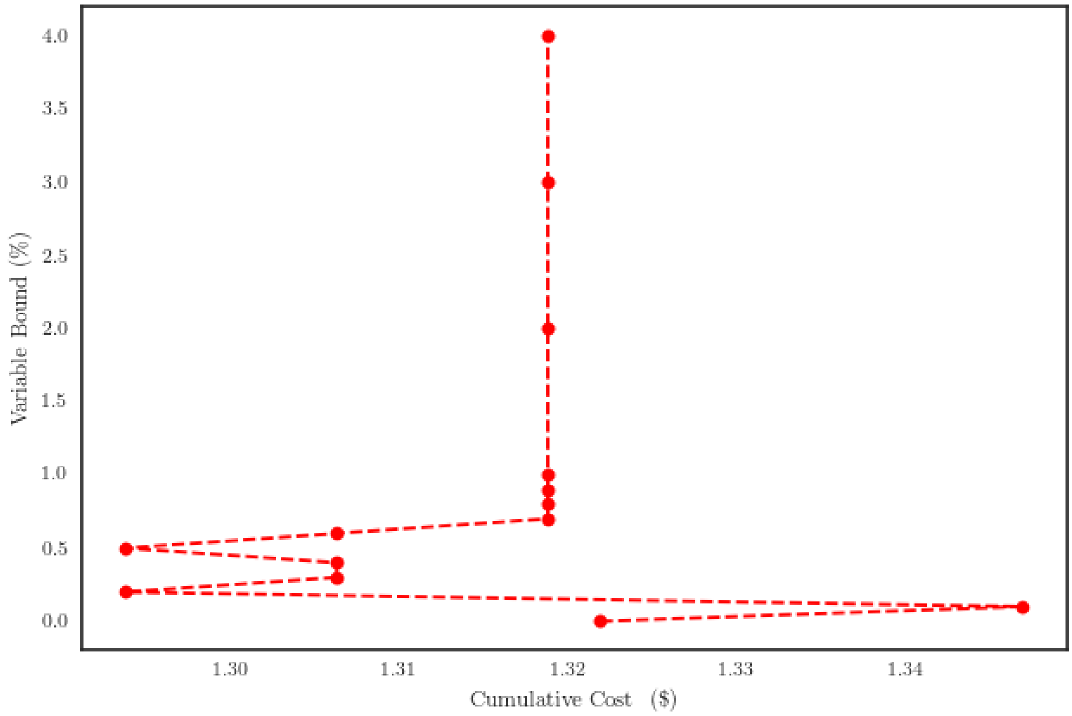

Figure 5.8: Balancing Cost for the Variable Bounding approach.

5.3 Weighted Sum Balancing

In this section we discuss our evaluation results against the general requirements stated in Chapter 1 (i.e., applicability, effectiveness, efficiency, usability) for the Weighted Sum Balancing approach. First and foremost, this approach is also applicable as it relies on minimal instrumentation and computational resources. In more detail, we observe that solutions are able to be captured in milliseconds, which render this approach suitable to meet the real time operation constraints.

Now, figures 5.9 - 5.12 illustrate the results of this approach for various weighting parameters. In more detail, figure 5.9 demonstrates the results of this approach, with respect to balancing thermal discomfort and cost for the weighting parameters within the range of 0.0 - 1.0. Furthermore, figures 5.10, 5.11 and 5.12 demonstrate the balancing of cost against discomfort, discomfort against the weighted sum λ parameter and cost against the weighted sum λ parameter for all weighting parameters respectively. As can be seen, clearly (especially in figure 5.10) the approach is able to capture Pareto optimal solutions effectively, is able to balance thermal discomfort and energy cost according to the human preferences and hence is effective. Furthermore, the approach captures Pareto solutions in the Pareto frontier, since we do not observe any majorly dominated solutions. Any small variation is attributed to planning approximations. In this context, the proposed approach meets the efficiency requirement stated in Chapter 1.

With respect to usability, the approach is able to capture a wide and evenly distributed range of solutions within the Pareto frontier, and as such is able to meet a wide range of user preferences, with respect to balancing energy cost and thermal discomfort. Furthermore, this is happening based on single parameter that is able to be populated by the user in an adaptive manner. For instance, if the user feels that the cost is high or that the discomfort is very low, the user can progressively reduce the λ parameter until their preferences are met. Conversely, if the discomfort is very high or the cost is very low the user can progressively increase the λ parameter until their preference is met, respectively. Hence, this approach also meets the usability criterion, as discussed in Chapter 1. In this context, this approach is more suitable for settings where a precise representation of the user preferences is required.

Importantly, with respect to Pareto efficiency, the efficiency of this approach is also supported theoretically, as the weighted sum scalarized function formulation of the multi objective optimization considers sufficient but non necessary condition for Pareto optimality (L. Zadeh 1963)[20]. In the next section, we provide a comprehensive discussion and a comparative analysis among the three approaches.

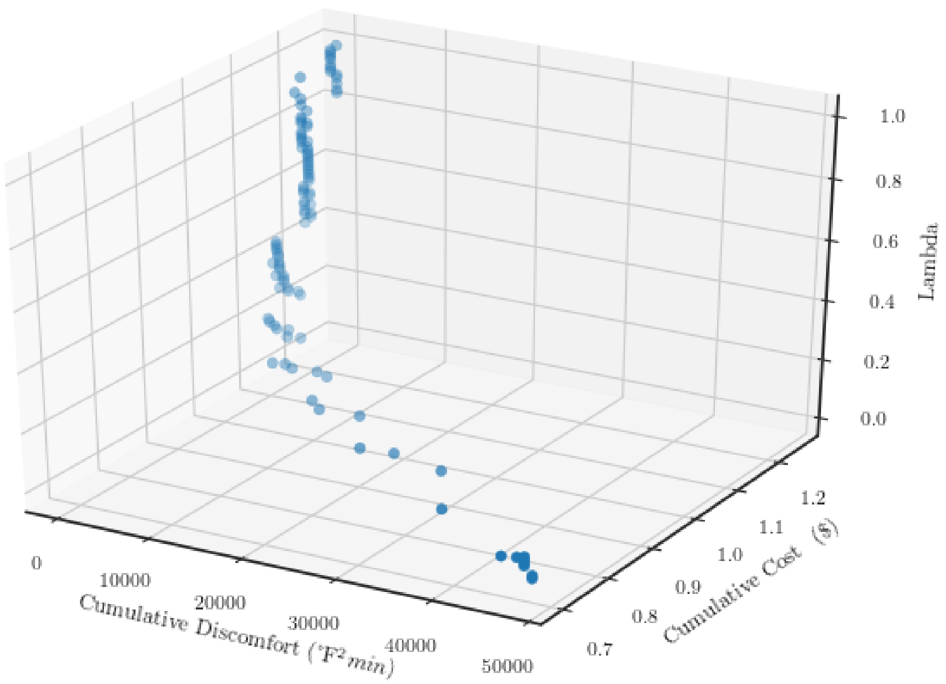

Figure 5.9: Aggregate Chart for the Weighted Sum Balancing approach.

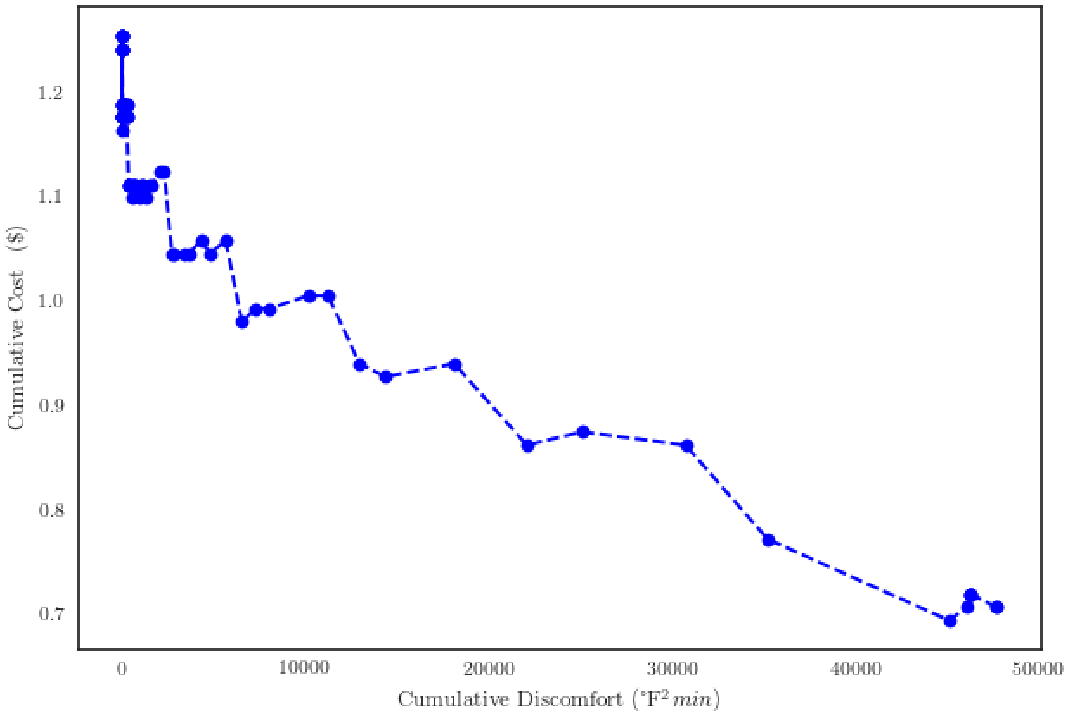

Figure 5.10: Balancing Cost and Discomfort for the Weighted Sum Balancing approach.

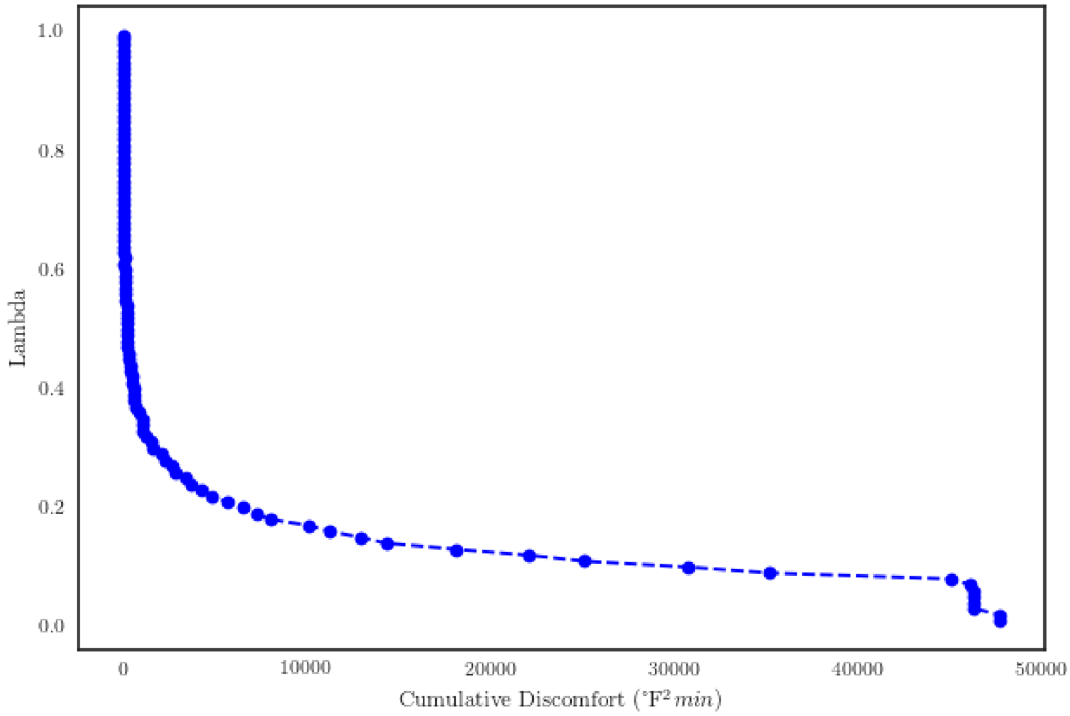

Figure 5.11: Balancing Discomfort for the Weighted Sum Balancing approach.

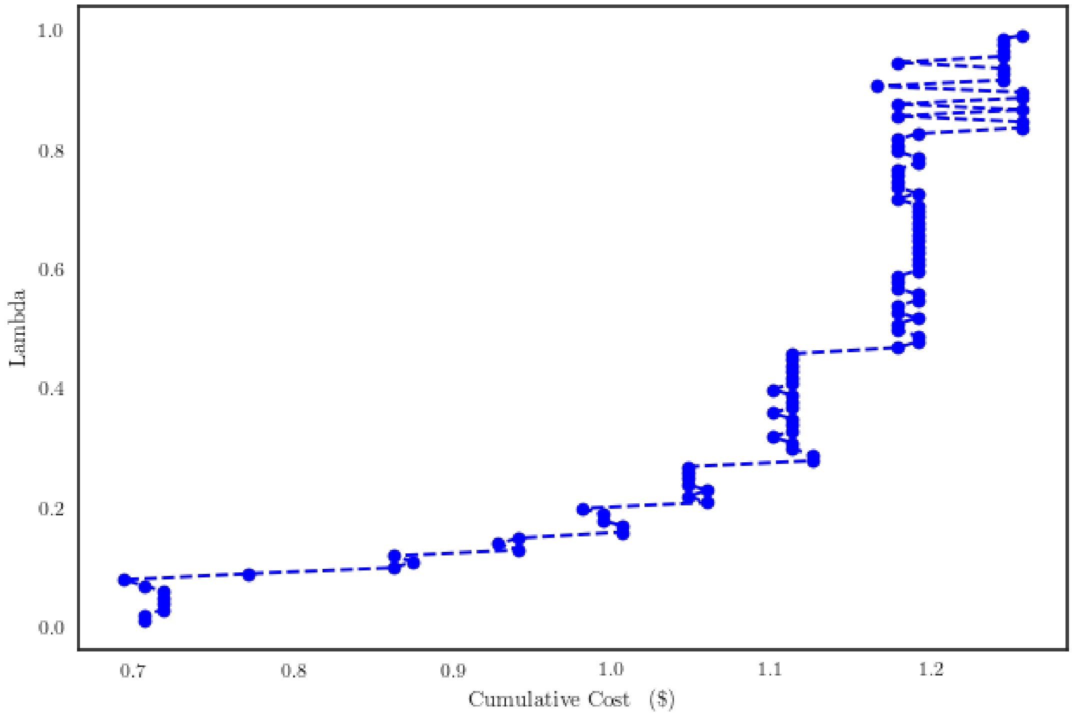

Figure 5.12: Balancing Cost for the Weighted Sum Balancing approach.

5.4 Main Results and Discussion

As discussed in sections 5.1, 5.2 and 5.3 all approaches, namely the Thresholding, Variable Bounding and Weighted Sum Balancing approaches, meet the effectiveness, efficiency, usability and applicability requirements. In more detail, all approaches are able to capture solutions in minimum computational time, that would allow the approaches to be applicable in settings with minimum computational resources and meet the real time operating requirements. Furthermore, all the approaches are effective, as they effectively balance the HVAC energy cost and the thermal discomfort of the residents, according to the user preferences. All approaches are Pareto efficient, as any dominated solutions observed are attributed only to the approximation of planning.

Nevertheless, the approaches vary considerably with respect to the usability criterion. In more detail, the Thresholding and Variable Bounding approaches seem to only capture solutions in a narrow range of Pareto optimal solutions within the Pareto frontier, and as such they are able to represent a small range of user preferences.

Hence, these approaches are suitable for cases where minimum thermal discomfort is required, with a minimum cost. In these settings, it is not even required for the population of the parameters for the Thresholding approach and the Variable Bounding approach respectively. Importantly, the Weighted Sum formulation is able to capture a wider and even distributed range of Pareto optimal solutions and thus it is suitable for settings where there is a need to capture a wide range of user preferences.

Our evaluation results are collectively reported in Figures 5.13, 5.14 and 5.15.

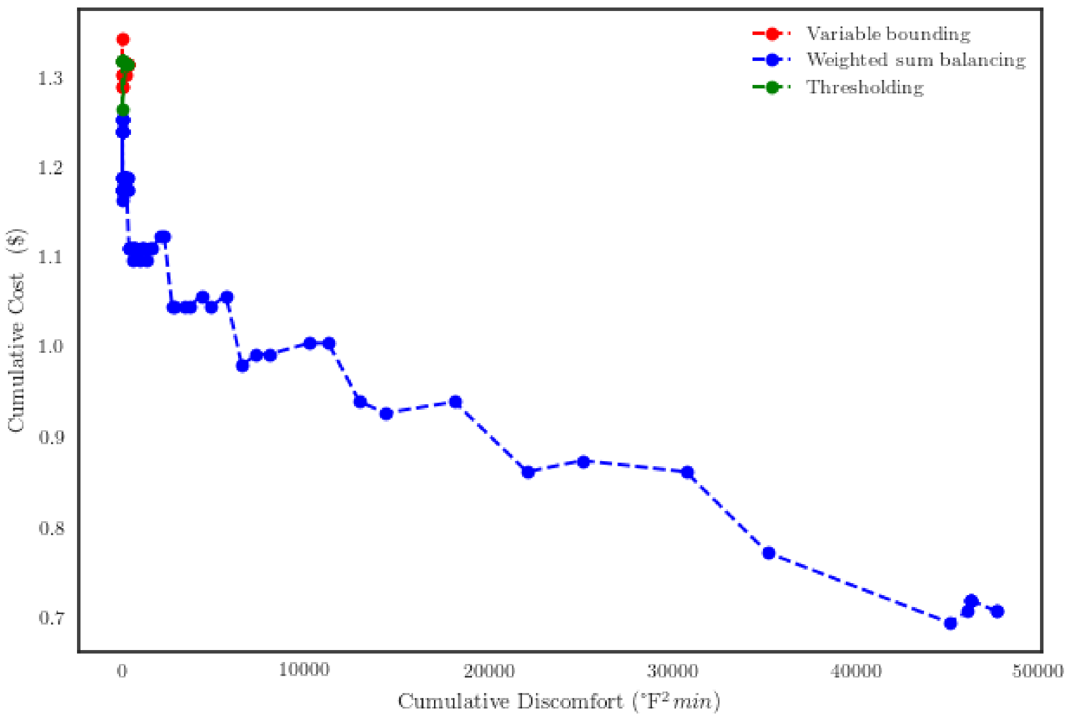

Figure 5.13: Balancing Cost and Discomfort for all methods.

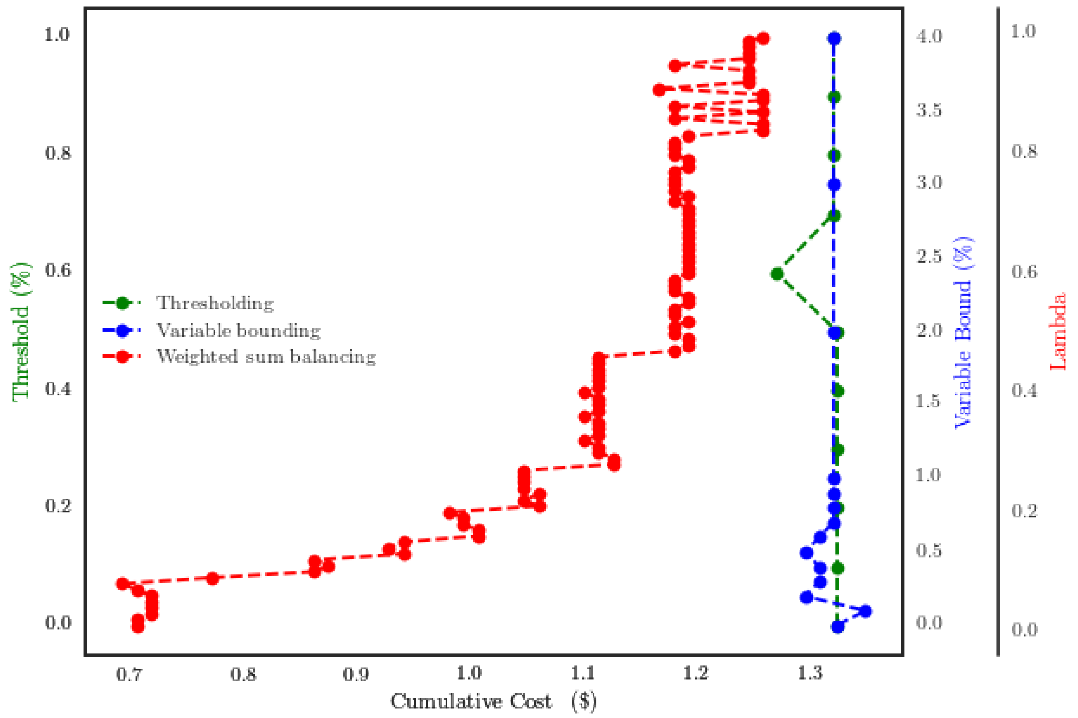

Figure 5.14: Balancing Cost for all methods.

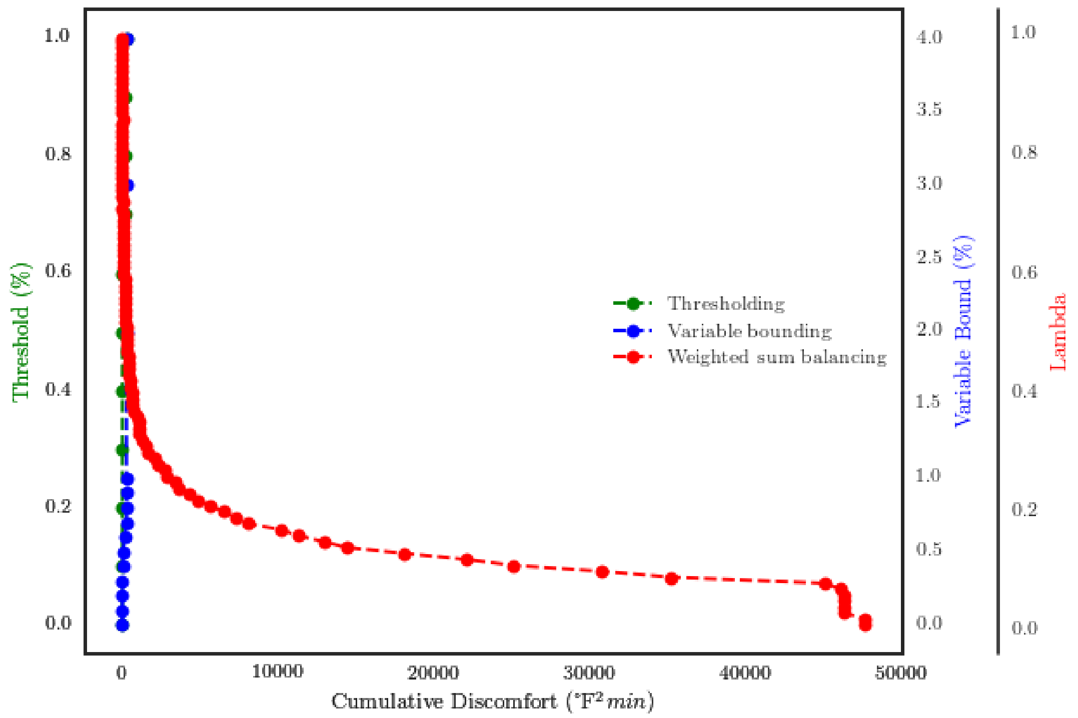

Figure 5.15: Balancing Discomfort for all methods.

Since both the Variable Bounding and the Thresholding approach produce solution points in a very narrow region, there is limited value in populating the balancing parameters with values other than the intuitive 0.5 for the threshold and the intuitive safety requirements for the bounds. This is further supported by the fact that small changes in the threshold and variable bound variables correspond to very small variations in the solution points captured, as can be seen in Figures 5.13 and 5.14 for both the Thresholding approach and the Variable Bounding approach.

Populating the Weighted Sum Balancing parameter is an important task. As discussed, Figure 5.15 depicts the relationship between some balancing parameter and discomfort, while Figure 5.14 shows the relationship between some balancing parameter and cost. Both cost and discomfort have a generally monotonous relationship with the balancing parameter which confirms that one can progressively increase or reduce the parameter until the user preferences are met in an adaptive manner (as also discussed in Panagopoulos et al. 2015).

Figures 5.16 to 5.20 illustrate the sample days that were simulated using the Weighted Sum Balancing approach, for concreteness purposes. As can be seen, the thermal comfort of the occupants is ignored for low λ parameter values, and the algorithm mostly respects the safety limits. When higher values are chosen for the λ parameter, the algorithm starts preheating the thermal zone, so that low thermal discomfort can be achieved when the occupants arrive, with minimum cost.

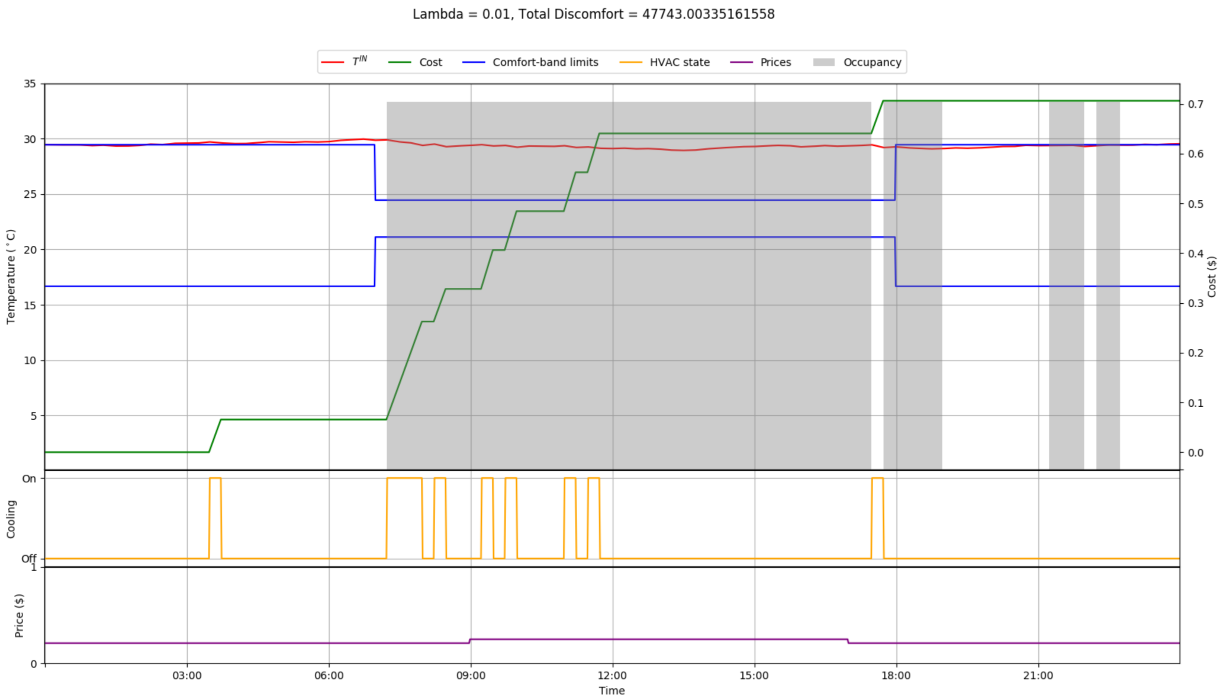

Figure 5.16: Sample day for $\lambda = 0.01$

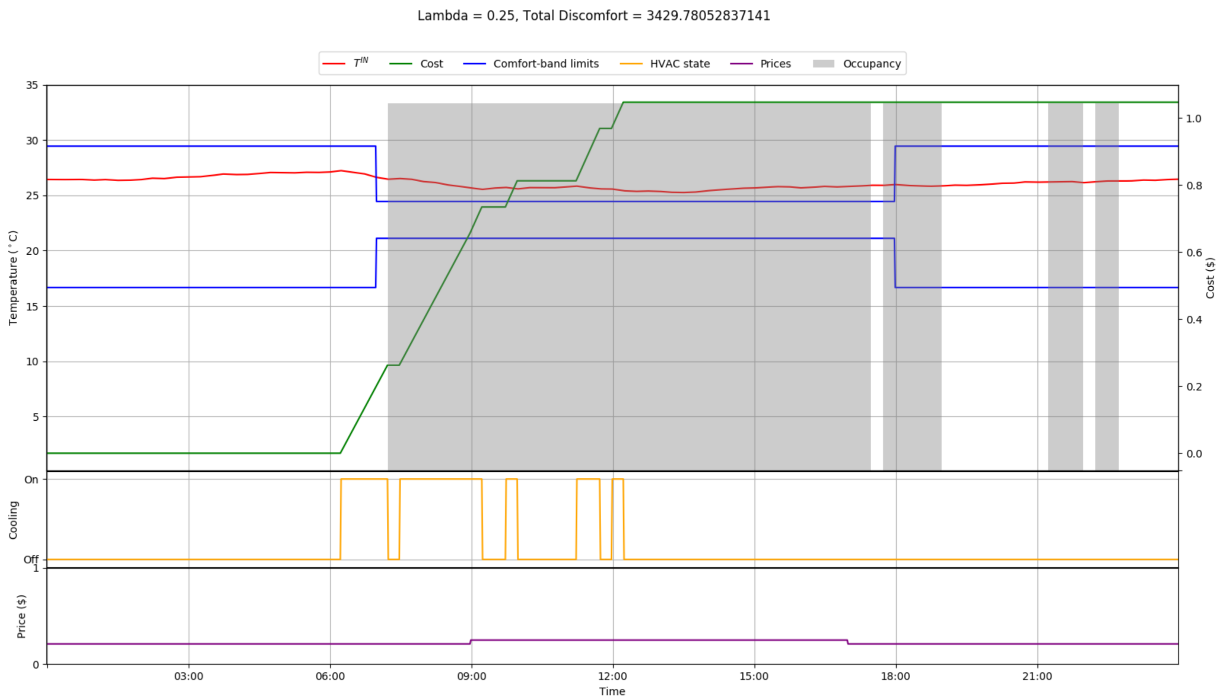

Figure 5.17: Sample day for $\lambda = 0.25$

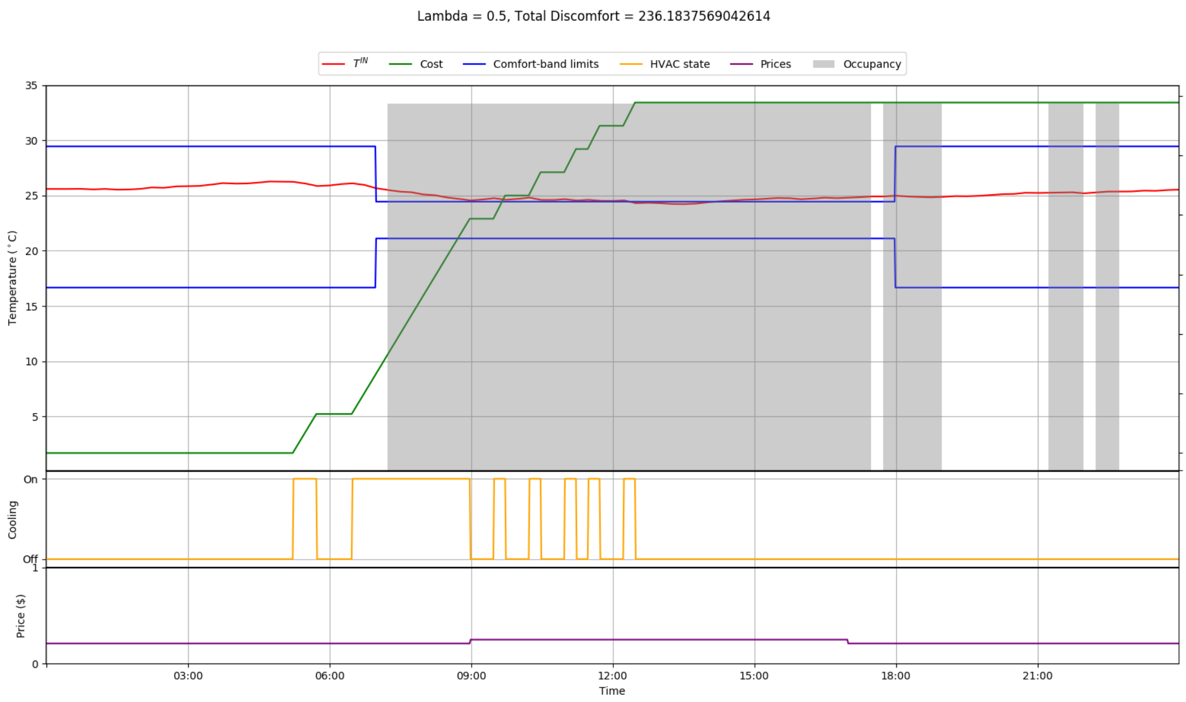

Figure 5.18: Sample day for $\lambda = 0.5$

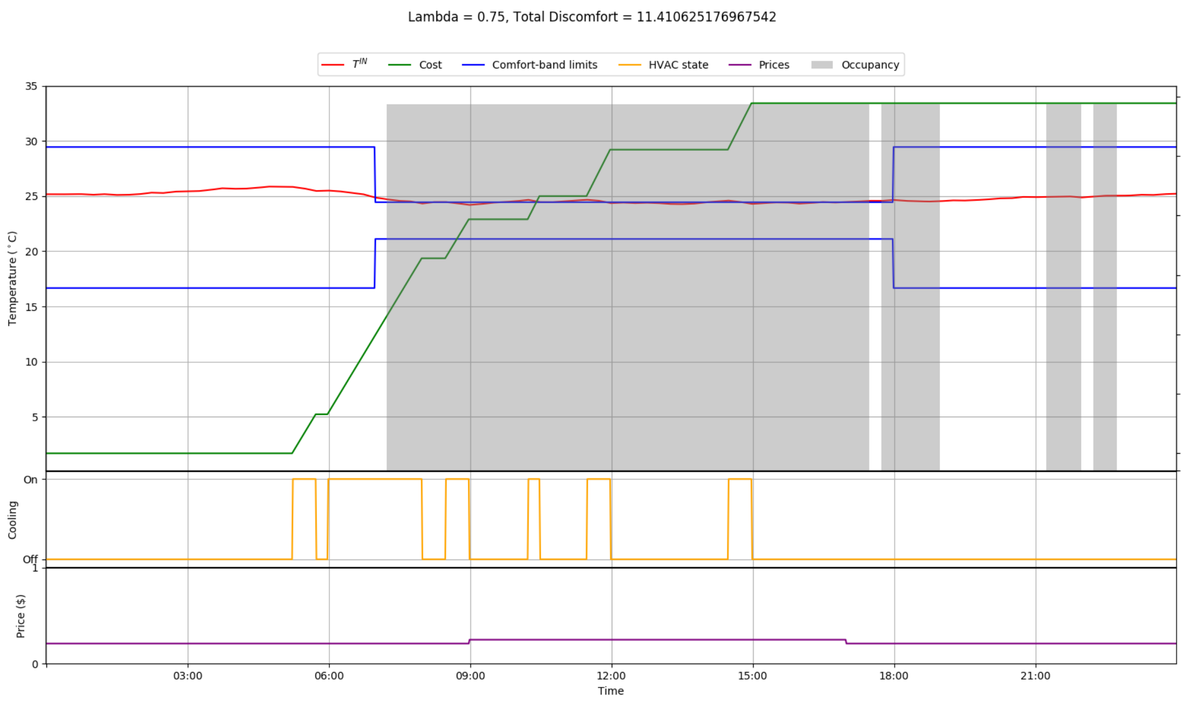

Figure 5.19: Sample day for $\lambda = 0.75$

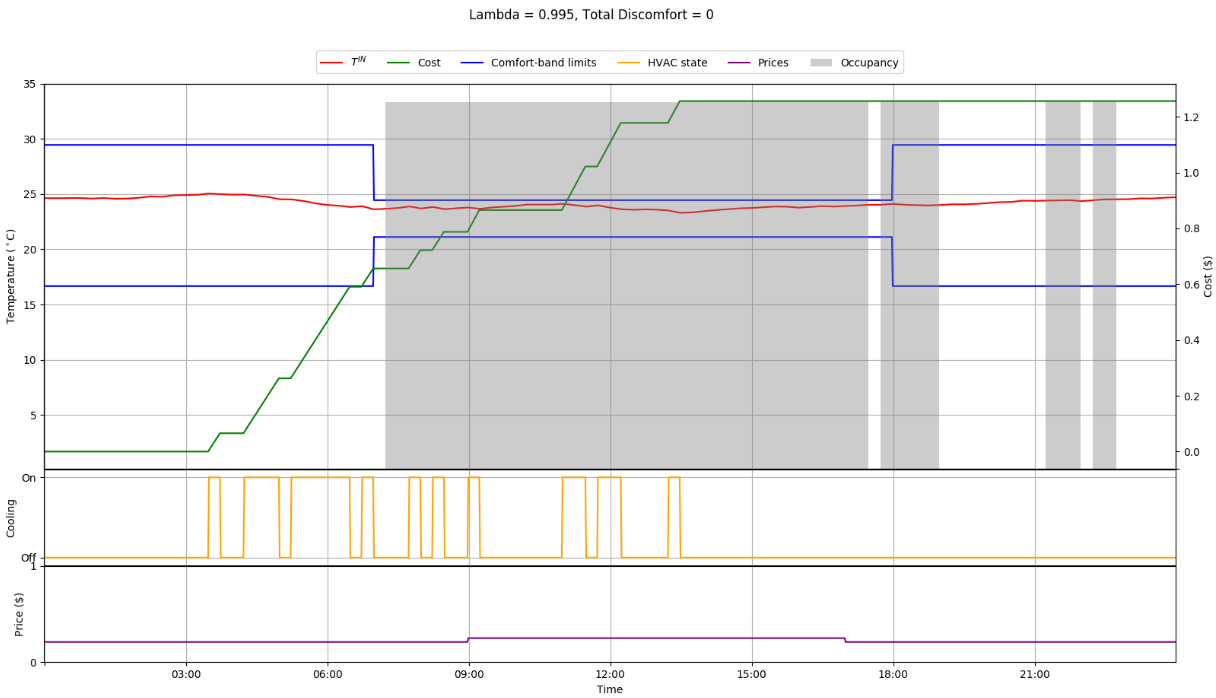

Figure 5.20: Sample day for $\lambda = 0.995$

To sum up, as discussed, although the Weighted Sum Balancing approach allows users to specify a wider range of balancing points, all approaches yield Pareto efficient solutions. As such, all methods are able to generate efficient solutions when a fixed operation is needed around the point of minimum discomfort. In addition, neither the Thresholding nor the Variable Bounding approach require users to populate any parameters.

Chapter 6

XBOS-DR and Real-world Trial Evaluation

As discussed in Chapter 1, this work aims to provide preliminary results to aid in the selection among various thermal discomfort and energy cost balancing approaches, to be used in a real trial evaluation as part of the XBOS-DR research endeavor. Based on the results discussed in Chapter 5, we selected the Weighted Sum Balancing approach. In the following paragraphs we discuss further lessons learned for our approach from the real-world trials.

In more detail, The XBOS-DR was a collaborative research project with partners from UC Berkeley, Technical University of Crete, Siemens and Quest, funded by the California Energy Commission with the aim to develop an operating systems for buildings that can support the seamless interconnection of Internet of Things (IOT) devices to enable the communication of such devices in different parts of a building. The aim of this building operating system is to support intelligent applications such as energy efficiency, demand side management, security diagnostics, etc.

In this context, the selection of a balancing technique between thermal comfort and HVAC cost serves as a preliminary evaluation in a simulation environment, before applying this technique to real world trials. The real world trials consider both a proof of concept for the operating system and also an evaluation for the technique itself in the real world. In order to evaluate the XBOS-DR operating system and the respective intelligent application ecosystem, including an intelligent thermostat HVAC control approach, we deployed 19 buildings across California as shown in the following table.

Site Name	Classification	No. of thermal zones
CSU Dominguez Hills	Business	8
Orinda Community Center	Multi-use assembly spaces	12
North Berkeley Senior Center	Senior center	3
The Local Butcher Shop	Mercantile	3
Avenal: Animal Shelter	Animal Shelter	13
Avenal: Movie Theatre	Assembly	13
Avenal: Veterans Hall	Senior Center	13
Avenal: Recreation Center	Community Center	13
Avenal: Public Works Department	Moderate Hazard Storage	13
Fire station 1, Hayward	Business	3
Fire station 8, Hayward	Business	3
Berkeley Corporation Yard	Business	3
Richmond Field Station, Bdg 190	Business	3
South Berkeley Senior Center	Senior center	3
Jesse Turner Fontana Community Center	Assembly	10
CIEE	Business	3
LBNL building 90C	Business	3
Word of Faith Christian Center	House of Worship and Accessory School Spaces	12
Orinda Library	Library	12

Table 6.1: Buildings selected for the project

These buildings consider small and average sized commercial buildings which are typically not employed with building automation systems, and hence XBOS-DR can fill this gap. The buildings range from butcher houses, that offer great application for thermostatically control loads for demand side management, to housing for the elderly that have more strict requirements with respect to the thermal comfort range. The experiments spanned over 2 summers in California, where appropriately designed demand side management signals called for reduction in the usage of energy.

In this context, three approaches for reducing energy cost have been evaluated. Among them, the Weighted Sum Balancing of energy operational cost and thermal discomfort which was identified by this work. The other two approaches considered the widening of the set-point temperatures and the business-as-usual strategy during a DSM event.

Preliminary results, demonstrated that both the Variable Bounding approach and the Weighted Sum Balancing approach were efficient in reducing the cost of energy within the aforementioned required range, according to the demand-side management signal. Nevertheless, the Weighted Sum Balancing approach (which was selected from this work) demonstrated slightly better performance with respect to energy savings in these events. In addition, the λ balancing parameter, which automatically adjusts the balancing of HVAC cost and thermal discomfort with respect to the price, has proven to be extremely user friendly.

In a nutshell, the preliminary results of this real world trial confirmed that the energy balancing approach of Weighted Sum Balancing can be supported by the XBOS-DR operating system, but also demonstrated efficiency and effectiveness when accounting for demand side management signals.

Imminent future work of the XBOS-DR and related project considers the incorporation of this balancing approach within a low-income thermostat appropriately designed for low income housing. This is supported by the relatively simple formulation of the approach, relying on a single scalarized function that renders the optimization simple to be solved, especially in devices with low computational power such as a low cost thermostat appropriate for low income houses. This line of research and progress, highlights the importance of our work of initially identifying this approach as a promising one with our simulation work.

Chapter 7

Conclusions and Future Work

In this work, we evaluated three strategies in stochastic occupancy settings, for the purpose of optimizing the HVAC control process and proposed a new approach that relies on variable bounding. Our analysis confirms that weighted sum formulation is superior with respect to the range of user preferences that is able to capture. Nevertheless, we show that all approaches evaluated are capable of capturing optimal solutions with minimum user input. As such, the preference of one approach over the others depends on the desired user input. In this context we also showed that in stochastic occupancy settings the setpoint temperature set by the user should consider the origin of the discomfort metric rather than a parameter to also balance discomfort and cost.

Future work includes experimentation with non linear expansion of the bounds in the variable bounding approach. Depending on the non-linear function utilized this approach can render variable bounding more risky or hesitant with respect to introducing discomfort in minimizing cost. Another goal is the incorporation of the humidity of a thermal zone as a part of the thermal discomfort modeling.

Additional future work includes the investigation of Peak Demand Charges within the context of our optimization schema, and the design of a distributed algorithm that would allow optimizing the HVAC control process of multiple thermal zones of a building, in a feasible manner. In more detail, multiple thermal zones typically consider large commercial buildings that come with advanced energy bills. Such advanced energy bills and billing schemas include peak demand charges. Peak demand charges require, apart from minimizing the total energy consumption within a specific billing period, to also trim the peaks, as high peaks lead to higher cost. As such, optimizing multiple zones in a way that synchronizes high energy demand processes to not operate at the same time can be crucial. Doing this in the context of the optimization MPC-based approach to respond to additional tariff-based DSM schemas, can be extremely valuable and considers a future work direction.

Bibliography

[1] Panagopoulos, A.A, Katsigiannis, M., Pritoni, M., Fierro, G., Lengyel, D., Peffer, T., Chalkiadakis, G,. Culler, D. Dealing with Expected Thermal Discomfort. ACEEE Summer Study on Energy Efficiency in Buildings Proceedings, 2018.

[2] Dounis A. I. and C. Caraiscos. Advanced control systems engineering for energy and comfort management in a building environment—a review. Renewable and Sustainable Energy Reviews, 13(6):1246–1261, 2009.

[3] Energy Information Administration (EIA). RESIDENTIAL ENERGY CONSUMPTION SURVEY (RECS), 2015. https://www.eia.gov/consumption/residential/index.php

[4] Fierro, G., D. Culler. An Extensible Building Operating System. Report UCB/EECS-2015-197 EECS Department, University of California, Berkeley, 2015. http://www2.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-197.html

[5] Gao P. X. and S. Keshav. 2013. Optimal personal comfort management using SPOT+. In Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings, pages 1–8. ACM, 2013.

[6] Lu J., T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. 2010. The smart thermostat: Using occupancy sensors to

save energy in homes. In Proc. of the 8th ACM Conf. on Embedded Networked Sensor Systems, pages 211–24. ACM, 2010.

[7] Panagopoulos, A.A., et al. 2015. AdaHeat: A general adaptive intelligent agent for domestic heating control. Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[8] Scott J., A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. 2011. Preheat: controlling home heating using occupancy prediction. In Proceedings of the 13th International Conference on Ubiquitous Computing, pages 281–290. ACM, 2011.

[9] Siroky, J., Oldewurtel, F., Cigler, J., and Privara, S. Experimental analysis of model predictive control for an energy efficient building heating system. Applied Energy, 88(9), 3079-3087, 2011.

[10] S. Prívara, J. Cigler, Z. Váňa, F. Oldewurtel, C. Sagerschnig, and E. Žáčeková. Building modeling as a crucial part for building predictive control. Energy and Buildings, 56(0):8–22, 2013.

[11] H. Huang, L. Chen, M. Mohammadzaheri, E. Hu, and MinleiChen. Multizone temperature prediction in a commercial building usingartificial neural network model. In Proceedings of the 10th IEEE International Conference on Control and Automation (ICCA 2013), pages 1896–1901, 2013.

[12] A. E. Ruano, E. M. Crispim, E. Z. E. Conceiçao, and M. M. J. R. Lúcio. Prediction of building's temperature using neural networks models. Energy and Buildings, 38(6): 682–694, 2006.

[13] ASHRAE Standard 55 "Thermal Environmental Conditions for Human Occupancy", (2013).

[14] Urieli D. and P. Stone. 2013. A learning agent for heat-pump thermostat control. In Proc. of the 2013 Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS), pages 1093–100. Int. Found. for Autonomous Agents and Multiagent Systems, 2013.

[15] W. Kleimingera, F. Matterna, and S. Santinib. Predicting household occupancy for smart heating control: A comparative performance analysis of STOA approaches. Technical report, ETH Zurich, 2013.

[16] Camacho E. F. and C. B. Alba. Model predictive control. Springer, 2013.

[17] Camacho, E.F and Bordons, C. Model predictive control, Springer-Verlag London Limited, 2004.

[18] REHVA, REHVA Journal 01/2012 - Energy-Efficient Passive House Using Thermal Mass to Achieve High Thermal Comfort., www.rehva.eu/rehva-journal/chapter/energy-efficient-passive-house-using-thermal-mass-to-achieve-high-thermal-comfort, 2012

[19] Marie Rousselot, Energy efficiency trends in buildings. Policy Brief, June 2018. https://www.odyssee-mure.eu/publications/policy-brief/buildings-energy-efficiency-trends.pdf

[20] L. Zadeh, Optimality and non-scalar-valued performance criteria, IEEE Transactions on Automatic Control, vol. 8, no. 1, pp. 59-60, January 1963,

[21] Osborne, Martin J. and Rubinstein, Ariel. A course in game theory. Cambridge, USA: The MIT Press, 1994.