Spring 1-1-2014

Stochastic Optimization of Building Control Systems for Mixed-Mode Buildings

Ryan Adams Tanner
University of Colorado at Boulder, ryan.tanner@colorado.edu

Stochastic Optimization of Building Control Systems for Mixed-Mode Buildings

by

Ryan Adams Tanner

B.S., University of Rhode Island, 2006

B.A., University of Rhode Island, 2007

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Civil, Environmental, and Architectural Engineering

2014

This thesis entitled:
Stochastic Optimization of Building Control Systems for Mixed-Mode Buildings
written by Ryan Adams Tanner
has been approved for the Department of Civil, Environmental, and Architectural Engineering

Gregor P. Henze, Ph.D.

Gail S. Brager, Ph.D.

Balaji Rajagopalan, Ph.D.

Michael J. Brandemuehl, Ph.D.

Clemens Felsmann, Ph.D.

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Tanner, Ryan Adams (Ph.D.)

Stochastic Optimization of Building Control Systems for Mixed-Mode Buildings

Thesis directed by Professor Gregor P. Henze, Ph.D.

Mixed mode (MM) buildings are a subset of low-energy buildings that employ both natural mechanical ventilation, often using manually operable windows for natural ventilation, along with other low-exergy cooling systems such as radiant cooling. This combination of systems has proven difficult to control in practice, in particular due to the potential for occupants to significantly impact building performance. Model predictive control (MPC) and rule extraction are promising methods for optimizing MM building systems in an offline setting, and for generating usable control rules that can be implemented in practice.

Simulation studies were performed to investigate the impact that occupant actions have on mixed mode buildings, and to improve the performance of natural ventilation controls in mixed mode buildings while accounting for uncertain occupant behavior. Results show that accounting for occupant behavior in building simulations provides useful insight into the robustness of different control strategies with respect to the impact of occupant actions. Two approaches to improving natural ventilation controls are applied to a physical building; the first seeks to improve on existing control logic by optimizing setpoints, while the second employs MPC and rule extraction to generate all new control logic. Each approach provides insight into potential flaws in existing logic and suggests revised logic that leads to better performance in the presence of occupant behavior.

In a final study, rule extraction is applied to optimal control datasets for multiple seasons and locations to develop control rules that approximate optimal controller performance. Converting state information to state-change information prior to applying rule extraction is shown to improve the performance of extracted rules, and it is shown that rules generated using data for a single season or location do not transfer well to other seasons or locations.

Dedication

For all of my teachers, instructors, coaches, and professors.

Acknowledgements

Many thanks to my advisor, Gregor Henze, for many thoughtful conversations, guidance, and for a truly limitless enthusiasm and drive to make this research fun, exciting, and successful. To my Ph.D. committee members and Robert Dodier for their advice and guidance throughout this project. To Shanti Pless, Jennifer Scheib and everyone at the RSF for making this work come to life in a real building. To Rose Grymes, Scott Poll, Rodney Martin and the entire NASA-SB staff for their support. To Peter May-Ostendorp and Chad Corbin, for building the MPC software and setting a strong foundation of work for me to build on. To Darren Robinson and Frederic Haldi for sharing code and insight into their models of occupant behavior.

Thank you to all of ASHRAE, ASHRAE-TC-1.4, and the RP-1597 PMS: Kimberly Barker, Chariti Young, and Michael Wetter, for endorsing, funding, and advising this research.

This work utilized the Janus supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado Boulder. The Janus supercomputer is a joint effort of the University of Colorado Boulder, the University of Colorado Denver and the National Center for Atmospheric Research. Thank you to Janus and the CU-Research Computing department, without whom this project would have taken 40 years in computing time alone.

To my friends, family, and fellow students, thanks for your endless support, understanding, conversation, distraction, and for keeping me sane.

Contents

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Motivation

As the population of humans on Earth grows, the world becomes more developed, and energy consumption increases, human activity has led to changes in the global climate. We, as a species, are changing our environment for the worse, and will continue to do so for the foreseeable future, feeding our addiction to fossil fuels to power our economies and our increasingly comfortable lifestyles. Renewable energy sources can help the situation by reducing carbon emissions from energy production, but the long term solution to the energy crisis plaguing our planet is not on the supply side of the energy equation, but on the demand side. We need to simply consume less energy, then the question of how to most efficiently or cleanly provide the energy need not even be asked; we can reduce both sides of the equation by demanding less. Less demanded equals less supplied.

On the demand side of the energy equation are four main sectors of energy consumption; transportation, industrial, residential, and commercial. Currently, the residential and commercial sectors combined account for 30% of energy consumption worldwide [95], and about 40% of energy consumption in the United States [96]. Within these sectors, heating, cooling, lighting, and appliance use are the main drivers of energy consumption. It is easy to save energy in any of these end-uses by simply turning the lights off, or turning the heat down - but realistically we need to figure out how to keep the lights on, keep the computers running, keep our buildings conditioned, and to do it all with less energy.

The work presented here focuses on a technique for conserving energy in buildings which marries together old technology and new in what are called mixed-mode (MM) buildings. Mixed mode buildings combine natural ventilation and mechanical ventilation (along with other low-energy sources of cooling such as ground source heat pumps) in a variety of ways in order to condition buildings more efficiently. MM buildings could fall in the commercial or residential sectors of energy consumption, but the primary focus here is on larger commercial buildings.

In the simplest design, a MM building can fully turn off its mechanical ventilation system and use natural ventilation whenever outdoor conditions are appropriate. By disengaging the mechanical ventilation system, fans, pumps, and potentially heating and cooling equipment are all turned off and the energy demand of a building is reduced. A method for further driving down the energy demand of a MM building is to capture free cooling by enabling the natural ventilation system to operate when outdoor conditions are cooler than indoor.

While the concept of a MM building can be simple and the operation of such a building is straightforward to describe, the reality is that MM buildings can be difficult to control in practice. A diverse mix of low-energy technologies are often combined in a single MM building, so there is no universal design and no standard control strategy that will work in every one. Additionally, an uncertain ingredient common to all MM buildings is the pool of occupants that use the building. Depending on what systems are available to the occupants in a MM building, e.g. operable windows, shades, fans, etc., the occupants may have a large potential to impact the performance of the building. Indeed, occupants can have an especially large impact on MM buildings that offer occupants access to operable windows, and may depend on occupant actions to keep the building comfortable.

In recent work, May-Ostendorp optimized MM building system controls in simulation, and used machine learning with simulation results to generate usable control rules that perform near optimally in simulation [63]. To cope with the complexities of MM buildings and system controls, May-Ostendorp used the simulation software EnergyPlus to model and simulate building performance, and a purpose-built model predictive control (MPC) environment to optimize control actions.

Optimizing building system controls is challenging due to the nonlinearity of thermal and fluid dynamics within building systems, the range of conditions a control system may need to work in, and the diverse time scales involved. EnergyPlus is a detailed simulation engine that is designed to cope with exactly these issues. The MPC environment allows one to optimize control actions over long time scales by breaking the problem down in time, optimizing one day at a time in sequence.

In order to leverage the results of controls that are optimized in simulation, May-Ostendorp used rule extraction to derive simple control relationships from optimal simulation results, and showed that the extracted control rules could work outside of simulation. The presented work expands on that of May-Ostendorp by including detailed stochastic models of occupant behavior in MM building simulation models, thus accounting for the impact of occupant actions on building performance, and by attempting to demonstrate the technology in a real building.

1.2 Research Objective and Questions

The central thesis of this work is that the added consideration of stochastic effects of human behavior in developing optimal building control strategies for MM buildings will lead to control strategies that are more robust to the impact of occupant behavior than strategies that are developed considering only mean-response occupant behavior, or none at all. The core questions we try to answer are listed below.

- Which aspects of occupant behavior have the greatest effect on energy consumption in MM buildings?

- How do the results of deterministic and stochastic MPC differ, and is the extra effort required for SMPC justified?

- Can we generalize extracted control rules to multiple climates or seasons?
- Do extracted control rules work in practice?

1.3 Organization and Summary

Having motivated and introduced the research presented here in Chapter 1, this section outlines the remainder of this document. Chapter 2 gives us a look at what the current state of the art is in mixed mode building research, predictive control in buildings, and occupant behavior in buildings, as well as some background on the use of machine learning algorithms in building control. Chapters 3, 4, and 5 give simple and detailed descriptions, respectively, of the methods applied in this research, and Chapters 6 through 9 show the results of the research.

The range of impacts that occupant behavior can have on a prototypical MM building are investigated in Chapter 6, with a focus on the impact of occupant actions on building heating, ventilating, and air conditioning (HVAC) energy consumption. The results show what one would expect - that occupant use of windows has the largest impact on HVAC energy consumption. The study also highlights the interaction between occupant use of building systems under different automatic control scenarios. A key finding is that the impact of occupant behavior on building performance is not consistent across all automatic building controls.

Leveraging the simulation and optimization environment built up in Chapters 5 and 6, Chapters 7, 8, and 9 follow two approaches to improving automatic window controls for a specific building. The first approach (bottom-up) takes an existing building control heuristic and explores opportunities to improve (to the extent possible) the performance of the existing controls by modifying setpoints in the existing logic. The second approach (top-down) follows a more complex two-step approach, where the first step is to optimize a long-term (i.e. multiple months) sequence of control actions, resulting in an optimal control decision time series, and the second step is to apply a machine learning algorithm to that optimal dataset to learn near-optimal control rules from the optimizer.

The majority of the work and simulation studies conducted throughout this project centered around the research support facility (RSF) building, which is a large net-zero energy building and part of the National Renewable Energy Laboratory (NREL) in Golden, CO. The final results presented in Chapter 10 consist of a description of the process followed to carry out a controls change in the RSF, and a presentation of the impact of the final controls implemented. Throughout the project, a simplified model of the RSF building was used in simulation studies to isolate potential faults in the window control logic, and to find opportunities to improve indoor comfort or reduce HVAC electrical energy consumption.

Interim conclusions and findings are given throughout each of the relevant Chapters, and the final conclusions are given in Chapter 11, along with some thoughts on opportunities for future work.

Chapter 2

Literature Review

2.1 Introduction

Five areas of literature relevant to this project are presented in the following,

- Mixed-mode buildings
- Occupant behavior in buildings and models of occupant behavior
- Model predictive control
- Stochastic processes and uncertainty in building simulation
- Rule extraction

2.2 Mixed Mode Buildings

2.2.1 Definition

Mixed-mode (MM) is a term used to identify a subset of high performance buildings (HPB) that use a combination of natural ventilation and mechanical ventilation systems to condition indoor air. MM buildings incorporate mechanical cooling and typically employ a variety of other low-energy cooling technologies, whereas strictly naturally ventilated buildings lack any mechanical cooling. MM buildings are designed to conserve energy by employing natural ventilation instead of mechanical ventilation whenever outdoor conditions allow, and extend the range of naturally ventilated buildings beyond temperate climates into more extreme climates where natural ventilation is possible in swing seasons, but mechanical cooling is necessary during the cooling season.

MM buildings often incorporate manually operable windows, blinds, or other systems that occupants can adjust to change their local environment and make themselves more comfortable. By enabling occupants to adjust local systems, MM buildings can take advantage of relaxed comfort standards such as the adaptive version of ASHRAE Standard 55 [4], or EN 15251 [68]. This benefits from both an energy and a comfort standpoint, since cooling setpoints can be elevated during warmer periods, and since occupant comfort is typically higher when occupants are able to adapt to their local environment [25].

The population of mixed-mode buildings in the US market is growing but as yet MM technologies are not specifically addressed in ASHRAE standards. While ASHRAE Standard 62.1 [5] does address natural ventilation, European and International standards adopt definitions for both hybrid ventilation and mixed-mode ventilation.. The International Energy Agency (IEA) and the European Committee for Standardization (CEN) define hybrid ventilation in Annex 35 (HybVent) [45] and EN 13779 [16] respectively. The Center for the Built Environment (CBE) at UC Berkeley, and the Chartered Institute of Building Services Engineers (CIBSE) in the UK provide similar definitions of mixed-mode building designs. For simplicity, the definition provided by the CBE will be used here and is as follows:

> Mixed-mode "refers to a hybrid approach to space conditioning that uses a combination of natural ventilation from operable windows (either manually or automatically controlled), and mechanical systems that provide air distribution and some form of cooling."[13]

The above definition is not very restrictive in terms of mechanical systems, and as such the variation in mixed-mode buildings is large. To account for this variation, a detailed classification scheme for mixed-mode buildings has been created. A summary of the classification scheme laid out by CIBSEs Application Manual 10 [20] is as follows. MM buildings are classified in two ways, first by the spatial distribution of ventilation strategies and second by their ventilation control topology, as in Figure 2.1 and in the following list.

1. **Zoned** In zoned MM buildings, different physical spaces are ventilated separately - each space is conditioned exclusively by either mechanical **or** natural methods.

2. **Complimentary** Complimentary systems are broken up based on temporal differences in control strategies into three sub-categories:

    1. **Concurrent** In concurrent systems both mechanical and natural ventilation can occur in the same space at the same time.

    2. **Changeover** In changeover systems, natural and mechanical ventilation do not occur at the same time and switching from one mode to another is controlled automatically.

    3. **Alternating** Alternating systems are similar to changeover systems but are usually switched between mechanical and natural ventilation for each season (e.g. natural ventilation is only active during swing seasons), rather than on an hourly or shorter time-scale.

3. **Contingent** The term contingent is used to describe any building that is a potential MM retrofit, or has the potential to utilize mixed-mode ventilation or cooling.

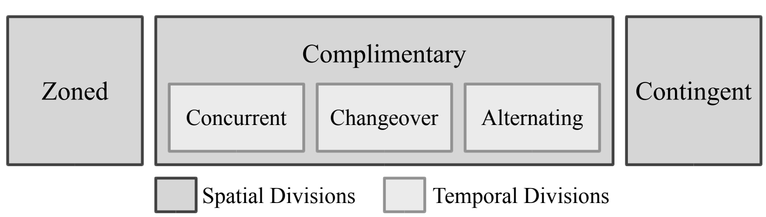

Figure 2.1: Spatial and temporal breakdown of MM designs.

2.3 Occupant Behavior in Buildings

2.3.1 Introduction

Building occupants can do a number of things to change the way a building performs. Occupant presence is itself an influence, resulting in heat gain, moisture gain, and generation of pollutants like CO_2. Occupant actions like adjusting windows, blinds, temperature setpoints, lights, and fans are some examples of the ways that occupants can alter the characteristics of a building zone, thereby changing the space conditioning loads, comfort criteria, and the performance of HVAC system. Occupant personal characteristics can also vary when occupants remove or add layers of clothing, eat or drink (which changes metabolic rates), or when occupants engage in more or less intense activities, e.g. sitting at a desk or working out in a gym.

Modern building performance simulation software calculates air, water, and energy flows within buildings, typically without accounting for realistic occupant-building interactions. Normally the impact of an occupant action (i.e. turning on a light or computer at work) is modeled as a regular schedule, i.e. lighting energy and associated heat gains are present from 8:00AM to 6:00PM on weekdays. Building occupants can influence all of these simulated building processes actively doing something, like turning a light on or off, or passively, by simply being present or absent. Depending on activity level, an occupant can generate more or less heat; in an office environment, workers use computers, printers and small kitchen appliances, all of which can increase noise and air pollution, generate heat, and consume electricity.

In attempting to simulate building performance more accurately, several questions emerge from examples like these. How much can occupants influence building performance? How, when, and why do occupants decide to take certain actions? What methods exist to accurately model occupant behavior in buildings? What level of detail is required to accurately simulate occupant behavior? Is a model of a single average occupant sufficient or should unique individuals be modeled explicitly? In the following review of pertinent literature, some of these questions are answered.

2.3.2 Occupant Behavior Models

In [68], Nicol et al. show that occupants take actions to keep their environment comfortable as ambient conditions change. This could mean opening a window or turning on a fan to increase airflow, or adjusting a thermostat to change the indoor temperature to a more comfortable value. When ambient lighting is below a comfort level, or daylight entering a space is uncomfortable, people are likely to adjust artificial lighting or activate shading controls to keep work plane illuminance at a desired level [52].

Intuition tells us that occupants take certain actions under certain generic conditions, e.g. when an occupant feels cold, he or she will take action to make him- or herself feel warmer, perhaps by increasing a thermostat setpoint or putting on a sweater, but it is difficult to accurately predict when an occupant will take action, and what action he or she will take. Accurate predictions are also difficult because different people might not perform the same action every time under the same conditions, (e.g. one person turns on lights when workplace illuminance drops to 200 Lux, while another turns on lights when it drops to 100 Lux). Similarly, a person might act the same way twice in a given situation, or he or she might act differently depending on season, time-of-day, or a number of other building-independent factors like mood, personality, or culture, confounding variables that are not available to use in predicting occupant actions.

Two primary methods for modeling occupant behavior have been developed; some models predict occupant actions (e.g. an occupant opens a window at time t), while others predict the state of some system which must have been acted upon by an occupant (e.g. a window is likely to be open at time t). The outputs of these models are inherently different, but the difference is subtle. To illustrate this subtle difference, consider a model of an occupant opening and closing windows. Within each type of model, the prediction of an occupant action is computed, but in the first type there are multiple output categories: (1) occupant opens a window, (2) occupant closes a window, or (3) occupant does nothing. In the second type, there are only two output categories: (1) the window is open, or (2) the window is closed. This subtle difference is somewhat trivial in building simulation, but it is more meaningful when building controls are considered, since building controls typically have to initiate an **action**, and cannot specify the **state** of a system explicitly.

The steps normally taken to develop a model of an occupant behavior (like those described later in this section) are as follows:

(1) Conduct a field study, recording ambient conditions and occupant actions.

(2) Find correlations between ambient conditions and occupant actions.

(3) Refine correlations into usable algorithms.

(4) Validate the model (function) by comparing its output to recorded data.

After completion of the field study, the second step is to create a probabilistic model of some behavior by correlating environmental factors (indoor and outdoor temperature, time of day, weather, etc.) with an **action** (opening a window, turning on a light, etc.) or with the **state** of a system determined by human actions (window is open, light is on, etc.). Frequently the model creators warn that site-specific data is probably leading to site-, building- or seasonally-dependent models, though models are kept as generic as possible to avoid these dependencies, and typically include additional 'tuning' parameters that are intended to adapt the models to different populations.

To illustrate this process an example of such a correlation between the status of windows and the indoor air temperature is given in Figure 2.2, adapted from [50].

In this example, the likelihood of a window being open is expressed as a function of indoor temperature. The probability p of a window being opened or closed is given by $p = at_0^2 + bt_0 + c$ where t_0 is the outdoor temperature, and a, b, and c are coefficients found for

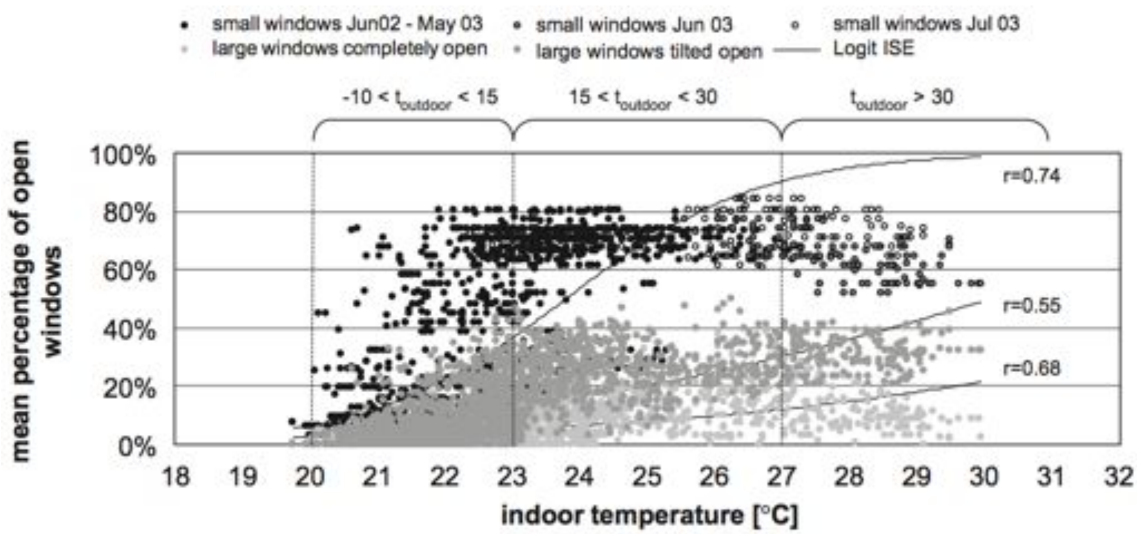

Figure 2.2: Correlation of open windows with indoor temperature. Adapted from Herkel et al. [50]

the probability function. Note that the coefficients a, b, and c, are not constant for a given person or day, and can be recalculated for a number of different situations. Further analysis of field data has shown that behaviors vary significantly at different times of day, typically occurring with a higher frequency around arrival or departure events. This dependence upon time of day or occupant transition is common among newer behavioral models, and is likely due to the occupant's reaction to changing from one environment to another, which can make the new environment seem more uncomfortable than it would otherwise (e.g., from walking to work in fresh cool air, and arriving in an indoor space that is warm and stuffy). That is, under the same indoor conditions, an occupant may be more likely to take an action when they have just arrived in the space than when they have been present in the space for a longer time. This temporal variation in behavior creates the need for multiple behavioral models, or multiple sets of coefficients to relate environmental conditions to occupant actions; i.e. one set for each state of presence (arriving, present, leaving). At this point, it may seem that we have potentially three probability functions per occupant, but studies have shown that considerably more complexity can be added to a model. For example, Herkel notes that different types of windows are acted upon differently, which leads to a different model for each type of window [50]. Along this same train of thought, it follows that a unique model could exist for different building façades - because offices on the west side of a building may be cooler in the morning than those on the east side - or different levels in a building - because windows on the ground floor are vulnerable to burglary and present a security breach when open. Further detail is often added to behavioral models to account for differences between occupants; unique models exist for active, medium, and passive lighting users [81], as well as for occupants who are more or less mobile (tend to leave their workspace more or less frequently) [75].

First generation behavioral models take the form of probability functions, simply relating the current state of a user-influenced system (e.g., an open window) to the current state of the environment (e.g., the indoor and outdoor temperature)[85]. Subsequent and newer models often take the form of Markov chains, which provide the probability of an action or state using both the current state of the environment and the prior state of the system being acted upon [37, 75, 103]. Markov chains tend to predict actions more accurately since they are dependent on the prior state of the system. Unless acted upon, things tend to stay in their current state; open windows are likely to stay open, lowered blinds are likely to stay lowered, and so on. Markov chains incorporate individual probabilities for each combination of states of a system; for a binary system (on/off), there will be four probabilities, the probability of changing state from on to off, or from off to on, and the two probabilities of remaining in either the on or off state.

It is important to note that each behavioral model described in this document is based on a unique set of data; some data sets are from a single office or building, while others draw on studies from multiple buildings in different climates. In each case, the available data could be leading to estimates of behavior that are unique to a particular climate, building, season, or culture.

2.3.2.1 Occupancy and Lighting Use

The purpose of a building is to provide a comfortable, safe, productive environment for people to work in, live in, play in - to exist in, thus building performance is driven by occupancy. For the purpose of estimating building performance, occupancy can be represented in many ways, from scheduled heat gains or scheduled appliance use, to aggregated values over the course of a given day, month or year for computational purposes. For example, when computing the cost of operating a desk lamp, we might simply assume that an occupant is present at the desk for 2000 hours per year, and that the occupant uses the lamp for 90% of that time, thus arriving at 1800 hours of lamp use per year. For computing more detailed metrics like building loads or air quality, it is important to know more precisely when occupants are in a building. When occupants are in a building during warm weather, they drive up cooling loads, and when present during cold weather, they drive down heating loads -

and a bulk metric like 2000 hours per year is useless without knowing **which** 2000 hours, so schedules are typically used in building energy simulation programs like EnergyPlus.

For the purpose of accurately modeling occupant behavior in buildings, modelers require even more detail about occupant presence and absence, not just knowing whether a person is present or absent, but knowing additional information on whether they have just arrived, or have been present for a while, whether they are about to leave - and whether their impending departure is for a short or a long absence. This extra level of granularity is required because every other occupant behavior (e.g. turning on a light) is contingent upon occupants being present to carry them out, and because these secondary behaviors are often linked to arrival and departure events.

The first attempts to develop accurate occupancy models other than schedules were driven initially by the lighting industry, when lights were almost entirely occupant controlled. Given that lighting was and continues to be one of the largest energy end uses in a building, and that it is often the easiest end-use to reduce, an understanding of occupant lighting use was useful in predicting and achieving savings. Given the tightly coupled history of the development of lighting-use and occupancy models, the two are presented together in this section. Perhaps the first, and certainly one of the most referenced occupant behavior models is the model developed in 1980 by Hunt [52], aimed at characterizing the interaction of occupants with lighting equipment. Hunt found the probability that a person would switch lights on or off based on the work-plane illuminance, which enabled better predictions of lighting use and energy consumption.

In 1995, Newsham and Reinhart [67] created a stochastic model of occupancy to increase accuracy in calculating lighting energy use; the model was based on typical occupancy patterns. Essentially, the new model added uncertainty to the times of arrival, departure, and breaks taken by occupants; it was assumed that occupants arrived and departed and took lunch within 15 minutes of a scheduled time, and the probability of taking morning and afternoon breaks was 50%. The two major outcomes of this work were a more realistic average occupancy profile and the program called Lightswitch. Lightswitch used observed occupancy data to produce realistic occupancy and lighting profiles [67]. At that time, the Lightswitch model was used to generate an average occupancy profile from multiple runs, and the average profile was then applied to every day in a simulation.

Seven years later, Reinhart refined and expanded upon the Lightswitch model to create Lightswitch 2002 [81]. This newer program added another degree of complexity and realism to the existing one by categorizing occupants as active, passive, or medium users (of lights and blinds). Lightswitch 2002 is used to determine annual energy savings using an annual occupancy schedule and annual work plane illuminance values to determine at each time during the year the actions an occupant will take on lighting or shading controls. Lightswitch 2002 enables users to quantify the range of potential energy savings from different levels of active or passive occupants due to their control of lights and blinds. The Lightswitch 2002 model has been implemented into a number of simulation and analysis tools (ESP-R, DAYSIM, and Lightswitch Wizard) and represents the industry standard for quantifying lighting energy use. While the development of occupant lighting use models has stagnated recently, interest in the development of occupancy models has grown.

Yamaguchi [102] proposed a model of occupancy in 2003 that represents the sequence of occupancy states (presence and absence, or 1 and 0) as a Markov chain. Driving this research was the requirement of an accurate estimate of occupancy to determine energy consumption from occupant appliance (computer) use, associated heat gains, and electricity use. Yamaguchi's model uses empirical data to define the distribution of arrival, departure, and lunch break times, and categorizes workers according to job-type. From these groups, energy-use profiles can be generated based on the percentage of workers in each job type, and used to calculate peak loads based on occupancy profiles.

In 2005, Wang [98] investigated one major shortcoming of existing occupancy predicting programs, namely that they did not address variations in occupancy throughout the day other than typical lunch and coffee breaks. Wang proposed "a probabilistic model to predict and simulate occupancy in single person offices", that would predict longer and shorter intervals of presence and absence throughout the day. Wang noted that vacancy and occupancy intervals are both exponentially distributed, and that the distributions of intervals vary throughout the day. Wang's final model accurately generates variations in daily occupancy (introducing periods of absence outside of normally scheduled lunch and coffee breaks), but assumes a normal distribution of arrival, departure, and break times, and that weekends are always unoccupied.

Noting several deficiencies in existing occupancy prediction models, namely long- and short- absences that occur in regular schedules, and non-traditional working schedules (i.e. working late, or working early) in 2008 Page [74] developed a more general model for generating occupancy profiles. Long and short vacations, business trips, and illness can cause a worker to be absent for longer periods ranging from partial days to weeks, and an impending deadline might make someone work outside of typical working hours or on weekends. These nuances of occupancy are addressed in Page's algorithm for generating occupancy profiles, which accurately predicts variations in occupancy throughout the day as well as throughout the year. The algorithm uses "an inhomogeneous Markov chain interrupted by occasional periods of long absence."

The number and duration of long absences are accomplished by generating probability distributions of the number and duration of absences using the inverse function method and a Poisson distribution with $\lambda = 3$ (this leads to a higher probability of three long absences per year). Subsequent to the definition of periods of absence for a given year, the remaining daily occupancy schedules are generated and a complete annual sequence of occupant presence and absence is realized.

Page's model was validated with two years of data from offices in the LESO-PB building at the EPFL in Lausanne, Switzerland, and can generate "a non-repeating time series of any length, including essential periods of long absence and otherwise reasonable movements to and from the zone resulting in an excellent estimate of the total time an occupant really spends within the zone simulated" [74].

In each of the three above listed methods for predicting occupant presence, empirical data is used to create probability distributions of occupant presence. In the models of Wang and Yamaguchi, the probability is time-independent, but in the Page algorithm, seven unique probability profiles are used - one for each day of the week, and the addition of long periods of absence for vacation adds a second layer of realism.

Finally, the Page algorithm also takes into account diversity among occupants by introducing the so-called mobility parameter, which describes how likely an occupant is to move into and out of a given zone. Specifically it is the ratio of sums of probabilities, where the numerator is the sum of probabilities of change in state, and the denominator is the sum of probabilities of **no** change in state. With relatively few inputs describing the probability of long absences, daily probability of presence profiles, and the parameter of mobility, the Page algorithm should be applicable to most any occupant (whether the model can be extended to more building types than office and residential is less clear). Chinnis et al. suggest additional tuning parameters that give the algorithm more degrees of freedom and make it more easily extensible to different building or space types [19].

2.3.3 Operable Windows

A summary of field studies conducted to investigate occupant behavior for window opening is presented here. These field studies were conducted in different countries, different climates, and had different setups and focus. They differed in observation periods (winter, summer, full year, short term, long term), office type (single occupancy offices, open office plans), as well as window type and facade design. Studies have typically found that the parameters most influential on occupant control are season, temperatures, time of day, and previous window state, but models differ significantly in both the predictor variables they choose, and their form.

Methods for quantifying the frequency and magnitude of occupant interactions with operable windows and shading controls have evolved from assumptions of behavior and simple heuristics to dynamic stochastic computations in the same way that methods for generating occupancy and lighting profiles were developed. Current models that deal with operable windows take the form of algorithms that use multiple probability distributions to determine when a given action (e.g., opening a window) will be taken, and the resulting system state. Often there are different probability distributions for different times of day, days of the week, or for different states of presence, i.e. arriving, present, or departing.

Rijal et al. [84] presented the Humphreys adaptive algorithm for simulating open windows in 2007. The Humphreys algorithm predicts the state of a window based on indoor and outdoor temperature using probability distributions drawn from field studies in 15 buildings in the UK between March 1996 and 1997; ten buildings were naturally ventilated and five were air conditioned. When implemented in the building simulation program ESP-r, the Humphreys algorithm predicts at each time-step whether occupants will open windows based on their comfort level (a function of indoor and outdoor temperature), and the resulting change in zone-airflow is used to determine the energy impact of occupant control of operable windows. The energy savings of different building designs was quantified by varying design parameters in simulations. Different design parameters were shown to influence occupant comfort, and thus occupant interaction with windows, and the energy-savings of different design options was found considering occupant behavior. In one example, it is suggested that improved building design reduces annual energy heating demand from 105 to 98 kWh/m2.

Bourgeois created a modified version of the Humphreys algorithm to explore the impact of defining users as active, medium, and passive with respect to window operation. Data from a study of one office building at the Universitè Laval in Quebec was used to generate the window opening model; the study spanned several months in the summer and fall of 2002 and the spring of 2003. This window opening model was only intended to determine if window operation models could benefit from detail on user characteristics in the same way that Lightswitch 2002 requires that occupants be defined as active, medium, or passive. The conclusion is that yes, indeed, models will benefit from defining window users as active and passive.

For implementation with building energy simulation software, Bourgeois [12] developed the sub-hourly occupancy-based control program (SHOCC) that works in parallel with ESP-r to enable behavioral models to be implemented dynamically within ESP-r (and potentially other building simulation programs). The Lightswitch 2002 model was implemented in SHOCC to predict occupant actions on lights and blinds dynamically in ESP-r simulations. Results of simulations revealed a potential for savings of 40% in primary energy expenditures between users that actively seek daylighting and users that rely on artificial lighting.

Yun [104] developed a stochastic model of occupant window-control and implemented it in ESP-r to demonstrate the influence of occupant actions on building ventilation. Results of simulations indicate that active window users effectively reduce the average summer indoor temperature by 2.6"C compared to passive users. The Yun model uses a Markov chain to determine the window state at each time step, based on the previous window state and probability distributions found in previous studies [103]. Monte Carlo methods are used in the simulation implementation to generate a distribution of results in accordance with the distribution of possible occupant actions during the simulation.

Herkel [50] created an algorithm to model window operation based on data from a year long study at the Fraunhofer Institute in Freiburg, Germany. This model predicts window status based on outdoor temperature and occupancy using an approach similar to Humphreys, with probability distributions of window control for arriving, present, and departing occupants. Herkel suggests that models could include further resolution with respect to season.

Haldi [37] created algorithms for window operation and the use of manual blinds using nearly 8 years of continuous data from a study at the LESO-PB building in Lausanne, Switzerland. His model for predicting window control uses the indoor and outdoor temperature, the prior window state, occupancy status and the presence of rain to determine the future (i.e. next time step) window state. The model takes the form of a Markov chain to determine when windows are opened or closed, and uses a Weibull distribution to compute the length of time that a window will remain open or closed. Haldi's algorithm for simulating user control of manually operated blinds follows the pattern found in the window algorithm; it uses outdoor horizontal illuminance, work plane illuminance, and occupancy status to predict when blinds will be acted upon and to what degree. This algorithm is unique in that it addresses a higher degree of uncertainty; not only is it unknown when occupants will act on blinds, but it is also unknown whether they will completely or partially raise or lower them; Haldis algorithm does predict both (it predicts the shading fraction conditioned on the raising or lowering action). This two-tiered or primary/secondary model arrangement provides another layer of detail and leads to more accurate representation of occupant actions.

2.3.4 Summary

Over the last decade, research on human behavior in buildings has led to the creation of algorithms that predict the probability of occupant actions, behaviors, or decisions based on data collected in field studies. Prominent among these are the Lightswitch 2002 algorithm for predicting lighting and blind use, the Haldi algorithms for predicting occupant use of blinds and operable windows, and the Page algorithm for predicting occupancy status. The majority of these algorithms have been implemented in popular building simulation software to demonstrate the influence that human behavior has on building performance, occupant comfort, and energy consumption. Results have shown that variations in human behaviors can lead to a 40% change in building energy consumption [12], and that using behavioral models in simulation can help to predict the energy cost or savings of different building designs. These models provide a sound foundation for further investigations of human behavior in buildings.

Given this discussion of occupant behavior modeling, and the recent developments over the last decade, it is no surprise that in November, 2013 the International Energy Agency (IEA) officially launched IEA Energy in Buildings and Communities (EBC) Annex 66,[1] which is a project focused entirely on occupant behavior simulation. The project is due to start in the fall of 2014, and will be complete in the spring of 2017. It seeks to leverage all of the work cited here, and is a collaboration between most of the authors cited in this section, as well as many others from around the globe. The goals of the Annex are to standardize the simulation methodology and models of occupant behavior, and to develop a consistent, usable framework for implementing occupant behavior models and understanding their impact on building performance through simulation.

2.4 Model Predictive Control in Buildings

2.4.1 Introduction

In the following, deterministic and stochastic model predictive control (MPC) are discussed with a focus on MPC in buildings. Initially hindered by complexity and lack of acceptance, these technologies have matured in the last thirty years to a point where they can be applied to building heating, ventilating, and air conditioning (HVAC) system control. Improvements in computational power have made offline MPC investigations easier, and MPC has shown strong potential through simulation studies and offline demonstrations for saving energy and increasing occupant comfort in buildings, though only a handful of real-time implementations exist. The main obstacles that must be overcome to realize MPC in buildings are (1) accurate system models, and (2) effective optimization strategies. These two components are imperative for MPC, and both present challenges to HVAC control because the controlled processes in buildings, i.e. airflow, water flow, and heat transfer are inherently nonlinear. Thus one might conclude that nonlinear system models and nonlinear optimization techniques are necessary for MPC in buildings. For real-time implementation

[1] http://www.annex66.org

of MPC, models must be simpler, which is achieved by linearizing non-linear processes and simplifying optimization problems to a point where they are computationally tractable.

2.4.2 MPC

MPC has its roots in the process industry, and excels where common PID controllers cannot achieve good control due to complex system dynamics and large time constants. Over the last three decades, computing and science have grown, enabling the application of MPC to nonlinear processes and systems with faster dynamics. MPC has found its way into the automotive, aerospace, and robotics industries, and in the last decade MPC has begun working for building control as well. MPC has proven to be robust by nature, even overly conservative in some cases because it guarantees (in theory) that the implemented control sequence will never cause the system to stray outside of its defined operating limits. MPC is also very effective at controlling highly nonlinear systems.

Before considering stochastic model predictive control (SMPC), it is important to understand deterministic model predictive control (MPC). Model Predictive Control is a method of process control, normally implemented in discrete time steps by repeating the following four steps.

(1) A process model is used to predict over a finite horizon the response and future states of the process given a sequence of control inputs.

(2) Optimal control sequences are computed to minimize some cost associated with the process.

(3) The first portion of the optimal control strategy is implemented.

(4) Time passes, and the system advances to the next time step.

As stated in the introduction, a model predictive controller relies heavily on two components: the system model and the optimizer, both of which present significant challenges.

Often, simple system models are adequate to achieve satisfactory control, as poor long-term predictions are discarded when the prediction horizon approaches the actual time of action, and more accurate near-term predictions are used instead. When system models are linear or convex, optimization is straightforward and follows well established mathematical guidelines for convex optimization, but typically MPC is applied to more complex systems, so practitioners transform nonlinear system models to linear form so that convex optimization techniques can be used.

Here we discuss SMPC in finer detail, but it is noteworthy to mention that MPC comes in a multitude of flavors and acronyms, including nonlinear (NMPC), open- and closed loop (O-MPC and C-MPC), robust (RMPC), receding horizon control, (RHC), adaptive predictive control (APC), and many others. The central concepts common to all variants of MPC are described above in the simplified four-step process, and throughout this document we use the terms MPC and SMPC to refer specifically to deterministic and stochastic model predictive control.

2.4.3 SMPC

Stochastic model predictive control extends deterministic MPC to more unstable processes, in particular those processes that are subject to stochastic variations in inputs and/or outputs. The defining feature of SMPC is that the controller assumes **a range** of possible system responses, and delivers a control strategy that is **very likely** to bring the system to a given state in spite of stochastic influences, while (deterministic) MPC assumes perfect knowledge of the system, and delivers a single control strategy (with a single predicted response) that may not be robust to the set of possible system dynamics. Note that stochastic MPC is not equivalent to robust control; robust control is theoretically robust to **all** possible disturbances, while stochastic MPC is only robust to the set of **likely** disturbances.

The most desirable and common treatment of SMPC for real-time implementation is with mathematically closed-form representations of the uncertainties. Often unknown disturbances, errors, or other criteria are assumed to be independent and identically distributed (IID) and Gaussian, which allows for simple computation of performance bounds. Such techniques provide some insight into what **might** be the range of possible errors or outputs from a model, but in reality these inputs often conform to non-normal distributions, and when the true distributions are unknown, the assumption of normality is no longer accurate and may lead to inaccurate results.

2.4.3.1 Addressing Uncertainty in MPC

Uncertainty can enter the model predictive control algorithm through several avenues:

- The **system model** is probably not perfectly accurate, so model predictions will be uncertain to some extent, even if inputs and disturbances are perfectly known.
- Uncontrolled **disturbances** that enter the system can be predicted but are also uncertain, and it follows that the effects of such disturbances are also uncertain as they propagate through the model.
- In real-time MPC implementation, **measurement** and **data transmission** errors or **time-delays** can skew the controller's interpretation of the true current state of the system, just as they would with conventional control methods (e.g. PID control).

In the presence of uncertain disturbances, a controller is faced with the difficult task of predicting an unknown future state and selecting a control strategy for the current time that will satisfy both present and future system constraints.

2.4.4 Approaches to MPC in Buildings

Predictive control of building systems is a relatively new field, with examples stretching back only a few decades. MPC efforts in buildings are focused on either low-level, local loop control, e.g. single-zone heating [58], or on high-level, supervisory control, typically leveraging thermal storage for cooling applications as in examples from Henze and Braun [48, 14, 49]. Some other examples of applications of MPC to buildings are included here, with extra attention given to examples that include mixed mode buildings, or stochastic effects.

May-Ostendorp [63] simulates deterministic MPC with a MM building modeled, where the building is modeled in EnergyPlus [91], and deterministic occupant behavior was considered, following the Humphrey's occupant window-opening behavior model [83]. Here, only the mean-response of occupant behavior was considered, so there was no variation from one simulation to the next in occupant behavior. May-Ostendorp's work and the work presented in this paper leverage the same framework developed by Corbin et al. for orchestrating simulated MPC with building models in EnergyPlus, as described in [24]. In Corbin's work, the development of the MPC environment is described in detail and a case study is used to demonstrate the feasibility of supervisory MPC in commercial buildings for cooling optimization when applied to the control of buildings with thermally activated building systems (TABS).

In [62], Mady demonstrates SMPC of indoor air temperature with a simplified single-zone model based on first principles, where occupant presence is the stochastic disturbance considered. Occupant presence is computed by an inhomogenous Markov model which provides a probability of presence for each hour. The SMPC controller uses the sensed indoor air temperature with the probability of occupancy to plan temperature setpoints for each minute to keep indoor air temperature constant.

In [30], Freire et al. demonstrate through simulation that MPC can be harnessed to provide superior occupant comfort while reducing energy consumption. A state space model of a single building zone is used, and five different cost functions are used with varying emphasis on comfort criteria and energy consumption. Comfort metrics included temperature and humidity bands, with a strong focus on the predicted mean vote (PMV) model. While a state space model of the zone was used by the controller, a whole-building simulation environment, PowerDomus, was used to simulate the MPC controller, which used a sequential quadratic programming optimization algorithm to discover optimal control signals.

Oldewurtel et al. present a methodology for approximating closed-loop constrained MPC in the presence of stochastic disturbances using affine disturbance feedback in [71]. To demonstrate the methodology, an example of building zone temperature control is given, in which stochastic disturbances include solar gains, outdoor air temperature, and occupancy (internal gains). Each disturbance is assumed to come from a predefined forecast, and is augmented by a random variable to induce variability; disturbances are assumed to be independent and identically distributed, and to come from a normal distribution. The example uses a state-space building model based on an RC (Resistance-Capacitance) nodal network.

Following on this work, Oldewurtel et al. applied SMPC to an RC network building model, in this case simulating uncertain disturbances with archived (real) weather forecasts and the corresponding forecast prediction error [72].

As part of the OptiControl Project[2] , Oldewurtel et al. [71] use affine disturbance feedback to approximate chance constraints on the MPC problem, as well as the MPC problem itself. In approximating the MPC problem, the vector of decision variables is broken down into adjustable ('wait and see') and non-adjustable ('here and now') decisions. The adjustable decisions are assumed to be affine functions of the uncertain disturbances that will occur in the future. Chance constraints are constraints on a system that must be met with a certain probability; in buildings for example, the indoor air temperature must fall within the comfort band with a high probability. Hard constraints are variables that have fixed maxima or minima, e.g. the output of a heating system; a 5000 BTU/h heater can not output more than 5000 BTU/h, or less than 0 BTU/h.

The MPC model addressed by Oldewurtel et al. in [71] is based on a linear time-invariant (LTI) system subjected to normally distributed disturbances. The assumption of normally distributed disturbances enables the chance constraints to be approximated by a hard constraint, turning the stochastic problem back into a deterministic one, as in [9]. The

[2] http://www.opticontrol.ethz.ch/index.html

resulting simplified problem is shown to be robust but less conservative than other robust MPC methods. Building on the methods described in [71], Oldewurtel et al. go on to demonstrate the simplified MPC controller in simulation.

An investigation of the potential of SMPC to save energy in buildings is provided in [72], in which Oldewurtel et al. present results from a parametric study of control techniques, building and HVAC variants, as well as climate variants. The simulated controller controlled HVAC air change rates, lighting, and blind positioning. Uncertainty in weather predictions was addressed by a Kalman filter comparing predictions to historical prediction and actual data. Three control schemes were compared: traditional rule-based control (RBC), certainty equivalent (CE) MPC, and SMPC. Certainty-equivalent MPC operates with the assumption that future disturbances materialize at their expected values, thereby neglecting uncertainty in weather or occupancy predictions; when these predictions prove false, CEMPC usually behaves poorly. In four out of six selected examples (each example representing a unique building or climate), SMPC resulted in fewer comfort violations than RBC, and in all six cases SMPC resulted in less building energy use. In a comparison of stochastic and certainty equivalent MPC, results showed that SMPC usually resulted in fewer comfort violations than CEMPC, and that SMPC generally resulted in less dramatic indoor temperature variations.

A recent example of MPC implementation in buildings is given in [90], where the heating system in a Czech Technical University (CTU) building in Prague, CZ was controlled via MPC. The building uses a thermally activated building system (TABS) for radiant heating, which was the system controlled via MPC; the heating supply water temperature was adjusted to minimize energy use and maintain occupant comfort. A two day planning horizon was found to be the best compromise between accuracy of weather predictions and allowed sufficient time for planning control strategies; the sampling rate of the implementation was 20 minutes. The building model was a data-driven RC network, and the optimization was achieved through a standard quadratic programming (QP) solver. Savings due to using MPC instead of PID control were documented between 15 and 23% depending on building construction and weather conditions. The authors point out that the methodology and implementation are neither simple nor trivial, but that the energy and associated cost savings will likely lead to widespread adoption of MPC techniques in the near future.

One recent example of MPC in the built environment is provided by Ma et al. in [59], and [60], in which thermal energy storage and building climate control are investigated. In [59], Ma et al demonstrate MPC of a campus-wide chilled water system at UC-Merced. Experimental results show reduced electricity cost and improved plant efficiency. System models are data-driven, and some input variables are implemented as look-up tables to reduce computation time and facilitate online implementation. Two major outcomes of the project are a 19% improvement in plant COP and a confirmation that existing system operation is near-optimal (the system managers were already operating the plant very well).

2.4.5 Summary

From this sampling of the literature relating to stochastic model predictive control in buildings, we note several major trends:

- System models are kept simple, leading to tractable, computationally efficient MPC formulations. With respect to building models, RC networks, state-space models, and data-driven models are the most popular. More detailed building models like those built in EnergyPlus, ESP-r, and TRNSYS are avoided due to the difficulty of implementation within an MPC scheme.

- Convex optimization is the normally chosen method due to its simplicity and familiarity; researchers are comfortable with convex optimization and build their models to fit within convex optimization schemes.

- Uncertainties affecting the system model are rarely (if ever) treated explicitly, and are usually approximated, transformed, or linearized to simplify the optimization problem.

- In recent building-related MPC experiments, uncertainties in occupancy [72], weather predictions [90], and deterministic occupant behavior [63] are addressed, stochastic occupant behavior is not.

Two main themes emerge as one looks at the goals of each building-system MPC effort. One theme is to use highly complex models of buildings and MPC to conduct **simulation studies**, which inform researchers about the potential for control improvements, and demonstrate the advantages of MPC. Another theme is to develop simplified models of buildings, processes, or stochastic influences that can be used in **real-time implementation**. In the latter theme, a common need is for computationally efficient models. The work presented here falls in line with the former theme, where complex, detailed, slower-running models of buildings and stochastic influences are used to give the most accurate simulation results and performance predictions possible.

An important note on the use of a more detailed simulation model (EnergyPlus), is that this type of simulation software is generally avoided in optimization studies because it lacks the ability to set states within the model or to control the numerical noise within the solver. The inability to set the state means that each simulation must be warmed up to the proper initial conditions, and the inability to control numerical noise within the solver can lead to rough cost surfaces, which in turn preclude the use of some optimization algorithms (those that require a smooth cost surface).

Given the complexity of MPC and a general lack of understanding or acceptance of MPC in the buildings industry, May-Ostendorp suggests in [63] that simpler, more understandable control rules will be accepted by building managers - and it has been shown that such control rules can be derived from the optimal datasets that result from simulated MPC of building systems.

In this project, using Monte Carlo sampling to investigate the impact of human behaviors leaves models of occupant behavior in their original stochastic form, rather than converting the impacts of occupant actions into min/max input disturbances, or assuming deterministic behavior.

2.5 Monte Carlo Stopping Criteria

Working with stochastic models can be challenging; care must be taken with the inputs (typically a random seed, and any number of weighting parameters or model coefficients), and with the outputs - which can be immense and difficult to analyze. For a deterministic annual hourly simulation, an analyst is faced with 8760 hourly values of each output variable, while for a stochastic simulation with 100 unique simulations, the analyst is faced with 100 realizations of each output, for each of the 8760 simulated hours. The analyst is then tasked with somehow turning this large dataset into something more tractable; aggregating 100 time series data sequences into one average, or perhaps showing the maximum of all 100 simulations, or using all of the data to show a distribution for each hour.

While the tasks of aggregating, summarizing, and visualizing stochastic results are a challenge compared to working with deterministic results, the more important questions are: "How many stochastic simulations are **enough**?", and "How do we define **enough**?".

In several uncertainty analyses of building simulation parameters, a typical range of 60-80 simulations is often cited [32, 57, 61] as an acceptable range, since the change in the distribution of a given output does not change substantially when more simulations are conducted. The question of defining the figure of merit for a suite of simulations depends entirely on the application. It could be an aggregate numeric like total energy consumption or average fan speed, a discrete metric such as the number of equipment cycles or the number of hours when comfort conditions are satisfactory, or a range, i.e. the temperature varied between T_{min} and T_{max}.

Since we are optimizing control sequences, and our optimizer requires a single number that can be used to rate the performance of a given control sequence, we use this same number, the objective function value, to determine if and when sufficient simulations have been conducted. A detailed discussion of the methodology employed to ensure convergence of the average of simulation results is given in Section 4.3.6.

2.6 Rule Extraction

In [63], May-Ostendorp introduces the notion of using machine learning techniques with time series building control data to **learn** near-optimal control strategies. Using simulation-derived optimal control time series similar to those generated in this work, May-Ostendorp demonstrates that rules based on generalized linear models (GLMs), classification and regression trees (CARTs), and adaptive boosting (AdaBoost) models all show promise as candidates for generating usable control rules for MM building systems. A number of optimizations were performed with four different MM building variants, and subsequent rule extraction was tested in different climates, buildings, and seasons to evaluate robustness of the rules. While the extracted rules in May-Ostendorp's work performed well in open-loop cross validation testing, some broke down in closed-loop testing, and a considerable amount of expert knowledge and manual tuning went into generating the rules. May-Ostendorp suggested that individual rules for different modes of operation, i.e. seasonal, weekday/weekend, occupied/unoccupied, etc., might be necessary, since rules trained on one dataset tend perform poorly under conditions that were not present in the training data.

In [26], Domahidi applies the AdaBoost algorithm to optimal control datasets for six different cases, which include different building configurations and climates. The control logic learned by the AdaBoost algorithm is shown to outperform heuristic controllers, approaching the performance of the optimal controller (called the 'Performance Bound') in all six cases. Domahidi also trained an AdaBoost model on the aggregate result of all six optimizations, and the rule trained with the aggregate dataset performed almost identically to the rules trained using individual data sets, when tested on the individual conditions. This result suggests that a larger dataset might enable training a single rule that works in multiple different buildings, or in multiple locations.

While May-Ostendorp applied rule extraction to both binary and continuous control inputs (window open/close, and temperature setpoints, respectively), Domahidi's work focuses on binary controls only, since binary decisions require mixed-integer programming optimization techniques, and are thus harder to implement in practice. Specifically, Domahidi looked at optimizing energy recovery and free cooling, since both of these systems are either active or inactive, and do not modulate across a range.

The methodology referred to later as top-down, is very similar to the methods used by May-Ostendorp and Domahidi, but is limited to the use of CART models since they result in the most readable control rules. Since clarity and simplicity in control rules are arguably the greatest barriers to acceptance by practicing building controls engineers and building operators, a primary goal throughout this project was to maximize the opportunity for implementation.

Chapter 3

Methodology

This chapter serves as a high-level overview of the more detailed discussion of methods explained in Chapters 4 and 5. Bearing in mind that mixed mode buildings present unique controls challenges due to the combinations of natural and mechanical ventilation systems that are used together in MM buildings, and due to the interaction of occupants with MM building systems, this project follows two paths towards optimizing automatic window controls for MM buildings. The two approaches are referred to as 'bottom-up' and 'top-down', where both begin with an existing building with some heuristic automatic window control logic, and both use simulation studies with a building energy model to find ways of improving building performance by changing the automatic window controls in simulation.

Both approaches begin with a building energy model, developed to accurately predict the impact of controls changes on the mechanical and natural ventilation systems in the actual building selected. In most of this work, the Research Support Facility, a large net-zero energy office building in Golden, CO, is the building in question. Since the building is large and complex, a full-scale model of the building proved too slow-running and cumbersome for use with optimization, so a simplified model of a typical cell of the building was used for controls investigations. Both methods also incorporate models of occupant behavior, which are coupled to the building energy simulation so that occupant actions are simulated in concert with the building simulation, leading to a distribution of results that show how occupant behavior impacts building performance. The final result of either approach is improved automatic window control logic that can be easily implemented into an existing building control system.

The bottom-up approach adopts the existing control logic, and makes adjustments to the setpoints and parameters within the existing logic to improve building performance. This is in contrast to the top-down approach, which begins with optimizing a sequence of control actions, then trying to find control logic that reproduces the same optimal sequence of control actions - or comes as close as possible to achieving the same level of building performance as the optimal sequence of control actions.

For the bottom-up approach, a simplified single-zone of an area of the RSF building is developed that includes all of the mechanical systems, natural ventilation, and controls in the primary office spaces of the building, as well as a model for simulating occupant use of operable windows. Since the same controls that exist in the building are also present in the model, it is possible to incrementally change the values of setpoints in the control logic, and observe the impact on building performance in simulation.

For the top-down approach, the same simplified model of the building is used as the objective function evaluator in a stochastic model predictive control optimization environment. In the control optimization, control actions are optimized one day at a time to arrive at multi-day (season-long) sequences of optimal control actions. Once a sequence of optimal control actions is available, it is used as the response variable, and combined with ambient conditions (predictor variables) to create a dataset of optimal control action time series. This dataset is then used by a machine learning algorithm to find causal relationships between ambient conditions and the control actions.

Occupant behavior is simulated either by coding probabilistic algorithms into Matlab and coupling them to EnergyPlus using the Building Controls Virtual Test Bed (BCVTB), or by coding the algorithms directly into the EnergyPlus input data file (IDF). The algorithms produce a probability of an occupant action at each time step, and comparison with a randomly generated number between zero and one determines whether the action occurs in simulation or not. Since this random comparator is different for each time step, and different for each simulation, multiple simulations are conducted to establish what the average impact of occupant behavior is on building performance. An investigation of how many simulations are required to ensure a reliable average impact is included in Section 4.3.6.

Early on in this research, the decision was made to use EnergyPlus to simulate building performance, and to use Monte Carlo simulations to evaluate the impact of occupant behavior on building performance. The combination of these two decisions, (1) to use a white-box, slow-running simulation model, and (2) Monte Carlo sampling, within an optimization scheme where thousands of simulations are required, led to a very computationally expensive process. This computational expense necessitated using the simplest EnergyPlus model possible to reduce runtime, and the smallest number of Monte Carlo samples possible to reduce the overall quantity of simulations.

Chapter 4

Control Approaches

4.1 Two Approaches

Recall that the major outcomes of this project are to determine the effects of occupant behavior on MM buildings, and to derive control strategies for MM buildings that account for building-occupant interactions. Two approaches to achieving these goals are illustrated in Figure 4.1 and described here; the specific tools and methods for each are discussed in the following sections.

4.2 Bottom Up: Optimization of Conventional Control Parameters

The first and simpler approach, the 'bottom-up' approach, entails parametric studies of setpoints in existing control sequences, outlined in the following procedure.

(1) Select an existing MM building.

(2) Construct a simplified energy model of the physical building.

(3) Incorporate the existing MM building controls in the energy model.

(4) Couple stochastic occupant behavior models to the building energy model.

(5) Choose setpoints from building control logic.

(6) Choose a set of values for each setpoint.

(7) For each unique combination of setpoint values, conduct a suite of simulations in order to capture the range of impacts of occupant behavior on building energy consumption and indoor comfort conditions.

(8) Select the set of control values that result in the best[1] combination of minimized energy consumption and maximized occupant comfort.

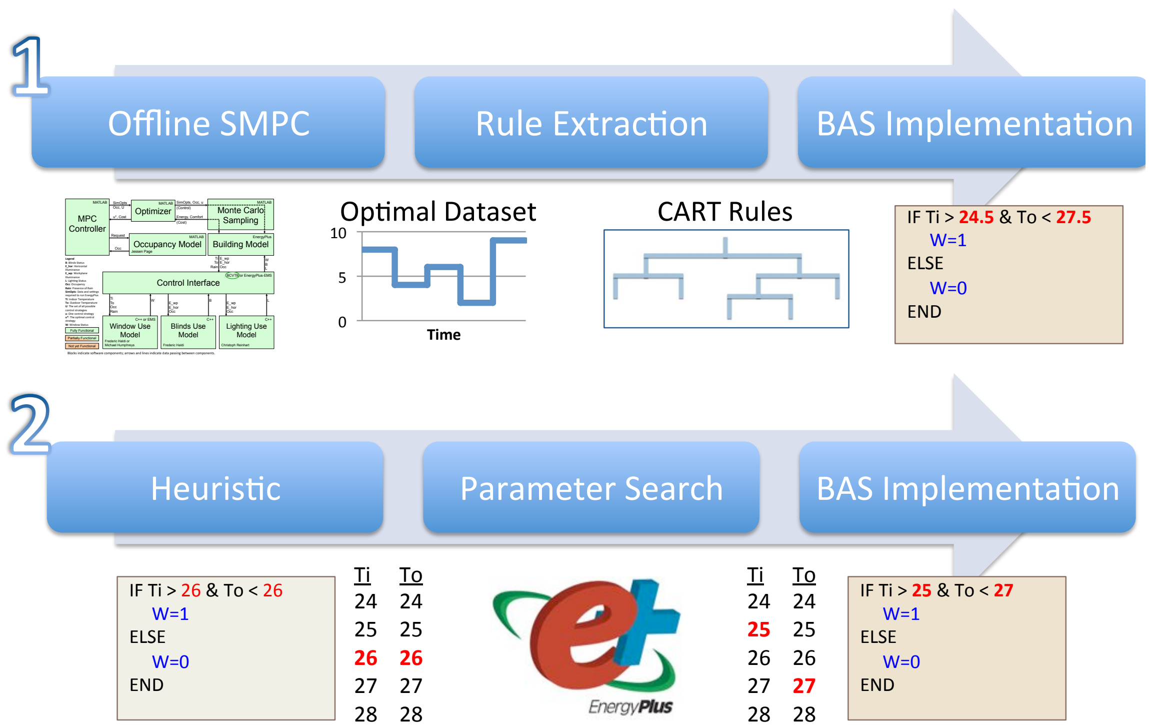

Figure 4.1: Two methodologies for improving MM building controls.

This procedure is both simple to implement and to understand, and should enable the tuning of existing MM controls to the best of their potential. The second approach, dubbed ‘top-down’, is considerably more complex, but may lead us to unconventional or unforeseen control strategies.

[1] The notion of what is **best** is for the investigator to decide; it could be wholly energy-focused, cost-focused, comfort-focused, a combination of multiple criteria, or some other objective entirely.

4.3 Top Down: Offline Optimization and Rule Extraction

The steps we follow in the second approach are nearly identical to the approach taken by May-Ostendorp et al. [63], with the only change being the inclusion of stochastic occupant behavior models.

(1) Select an existing MM building.

(2) Construct a simplified energy model of the physical building.

(3) Incorporate the existing MM building controls in the energy model.

(4) Conduct offline SMPC studies with the building energy model, arriving at a dataset that includes a sequence of optimal setpoints.

(5) Use rule-extraction to find relationships (control rules) between measurable ambient conditions and optimal setpoints.

(6) Tune and validate extracted control rules.

The key difference between the methods in [63], and those described here is in the treatment of occupant behavior. Here, instead of simulating the building energy model once per objective-function-call, we simulate the model multiple times per objective-function-call. It is shown in Section 4.3.6 that we simulate enough times to arrive at a representative expected value of the objective function (building energy consumption, occupant comfort, etc.).

In the next section, software environments and methods used to undertake the above two methodologies are described.

4.3.1 Supervisory Control

This work does not address local-loop or terminal unit controls, rather it focuses on whole-building supervisory control, and seeks to determine control logic that will take advantage of whole-building features such as thermal mass or natural ventilation to improve comfort conditions and decrease energy consumption. The outcome of a given study that follows this methodology is a rich dataset which includes simulation results for a building being controlled by two distinct strategies: a heuristic control strategy, and an optimal control strategy. All results include distributions which show the impact of occupant behavior on building performance.

The means by which we control the entire building in simulation is via setpoints, which could be temperatures, flow rates, or some other threshold values in a control sequence. In the example provided in the results section, we optimize automatic window control signals, which can take binary values of either 0 (close) or 1 (open). As we look at longer time-scales, i.e. building operation for many days or months, optimizing a setpoint for each hour becomes increasingly computationally expensive. To address this expense, we break up the problem in time with model predictive control (MPC), advancing through each day in a month, and optimizing each day individually instead of the whole month at once.

4.3.2 Organization

This methodology section is broken up into four components. In the first portion, we address the implementation issues associated with EnergyPlus, which are also addressed in detail by [24]. In particular, we discuss how the objective function is broken up into different temporal components, explaining how each component applies to a different time-horizon required for MPC with EnergyPlus.

The next section continues with a discussion of the objective function, now delving into the specifics of how we provide a single metric for an entire simulation. Where the first section discusses the breakdown of the objective function into different values for different periods of time, the second section discusses how the objective function is broken down into two components: one for energy consumption, and one for comfort conditions.

In the third portion of the methodology section, we discuss the details of the stochastic phenomena that impact the sequences of events in each simulation. This section highlights the differences between stochastic and deterministic MPC, and what care needs to be taken when conducting suites of simulations instead of one for each unique case. Here we introduce several different simulation cases, the case that uses default controls (DEF), the case that uses optimal controls (OPT) and the cases that the optimizer chooses from, called candidates (CAND).

The fourth element of the methodology is a discussion of the number of stochastic simulations required to provide a satisfactory distribution of results, or the Monte Carlo number. Often, this number is chosen a priori as some round number such as 50, 100, or 1000, depending on the goals of the researcher, and on computational constraints. Here we implement a variable convergence criteria that ensures we have conducted enough simulations, and that we do not needlessly simulate too many simulations. In practice, we ensure that at least 50 simulations, and no more than 400 simulations are conducted.

For clarity, a nomenclature section with terms used in this chapter is included here.

4.3.3 SMPC in Simulation

In most MPC applications, simplified (typically linear) dynamic models are used as they allow for many simulations in a short time (i.e. thousands per second). In the application presented here, a more complex model of an entire building, modeled in EnergyPlus, is used which presents several major issues. The first is that initial and final state variables cannot be directly prescribed, the second is that there are more states in an EnergyPlus model than in simpler models, and the third is simulation time. Simulation time is minimized wherever possible throughout the approach, and the treatment of simulation start and end states is defined here. We begin by stating the MPC problem mathematically, then describe the time-horizons pictured in Figure 4.2 and explain why each is necessary in this approach.

Model predictive control is a method of dynamic optimization in which multiple control decisions for a dynamic system are optimized sequentially. For a given point in time t, control

Table 4.1: Nomenclature.

Symbol	Definition
C	The value of the objective function.
P	The portion of the objective function computed over the planning horizon.
T	The portion of the objective function computed over the termination horizon.
R	The cost of energy.
M, M'	The cost of occupant-weighted comfort conditions.
N	The number of execution horizons in the planning horizon.
K	The number of execution horizons in the termination horizon.
$\vec{u}$	The vector of control decisions at a given point in time.
$\vec{x}$	The vector of system states at a given point in time.
B_W	The bandwidth inside which the moving average must fall to ensure convergence.
B_L	The number of consecutive simulations required to reach convergence.
t	A time index.
n, k, i	A simulation index.
DEF, d	Refers to the default case.
$CAND, c$	Refers to a candidate case considered during an optimization.
OPT	Refers to the optimal case.
N_{MC}	The number of Monte Carlo simulations.
S	A simulation.
F	The final candidate in an optimization.
W	The set of possible disturbances.
W_R	Weighting coefficient applied to the cost of energy.
$W_M, W_{M'}$	Weighting coefficients applied to the cost of comfort.
U	The set of possible control actions.
$\forall$	For all.
$\overset{!}{<}$	Must be less than.
$\lceil \cdot \rceil$	Ceiling.
$(\cdot)^+$	Positive values only.

decisions are optimized as in Equation 4.1

$$\min_{\vec{u}\in U} \sum_{t=t_0}^{N} C(\vec{x},\vec{u})\,\forall\, W \tag{4.1}$$

where $C(\vec{x},\vec{u})$ is an objective function describing the cost of operating the system with states $\vec{x}$ and a vector of control decisions $\vec{u}$.

In stochastic MPC (SMPC), the state at any future time is uncertain due to stochastic disturbances in the system, and the new formulation is

$$\min_{\vec{u}\in U} \sum_{t=t_0}^{N} E[C(\vec{x},\vec{u})\,\forall\, W] \tag{4.2}$$

where $E(C)$ is the expected value of the cost function, and W is the set of possible disturbances. Once the optimization for a given point in time is complete, the first element of the control sequence $\vec{u}(t = t_0)$ is implemented, and the system steps forward in time (typically in discrete time steps) to $t = t_1$. This is referred to here as one execution, and the time between optimizations we refer to as the **execution horizon**, as shown in Figure 4.2. At each time step, the objective function given in Eq. 4.3 is used to evaluate performance for a given simulation:

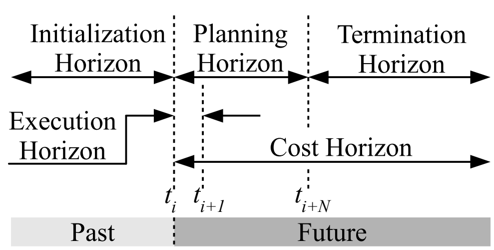

Figure 4.2: Time-horizon nomenclature.

$$C(t) = \sum_{t=t_0}^{N-1} P(\vec{x}(t),\vec{u}(t)) + T(\vec{x}(t_N)) \tag{4.3}$$

where $P(\vec{x},\vec{u})$ is a cost function for the planning horizon and $T(\vec{x})$ is the cost associated with leaving the system in the terminal state $\vec{x}(t_N)$.

In a simple system with only a handful of states, it is tractable to assign penalty terms to each state, however in a complex model and in a general sense, it is difficult to assign a penalty to every possible state. For example, if a problem involves the surface and air temperatures in a zone, combined with flow rates and water temperatures in a heating system, should the penalty be high when flow rates are high, or when temperatures are high at the end of a simulation? To address this issue, we assign a cost to the whole building performance during the entire termination **horizon** (see Figure 4.2), rather than to the terminal **state** at time t_N. This new formulation of the objective function is given in Equation 4.4.

$$C(t) = \sum_{t=t_0}^{N} P(\vec{x}(t), \vec{u}(t)) + \sum_{t=N+1}^{N+K} T(\vec{x}(t), \vec{u}_T(t)) \tag{4.4}$$

where u_T is the control sequence applied during the termination horizon, and K is the number of time steps included in the termination horizon. This leads to a new question: what controls should be applied during the termination horizon? Typically in MPC the planning horizon is the only time period considered by the optimizer. Since optimizing more than 24 hourly decisions is computationally intractable for the EnergyPlus models used in this work, we cannot reasonably optimize 48 hours at once, yet it is still important that we consider the impact of decisions made during one day on building performance for the following day, particularly for buildings with high thermal mass. So, while we can optimize a single 24-hr period in detail (the planning horizon), we must assume some heuristic control is applied during the termination horizon for the purpose of computing the terminal cost. In order to account for long-term effects of decisions chosen for the planning horizon, we assume that the building is operated with default control logic during the termination horizon, so that the results of candidate simulations can be compared with results of default simulations, and the only differences in cost can be attributed to differences in controls during the planning horizon.

In nominal MPC, Figure 4.2 can be simplified to the planning and execution horizons, with initial states $\vec{x}(t_0)$ set explicitly, and terminal states $\vec{x}(t_N)$ computed at each time step, however with EnergyPlus as the simulation engine, the full set of horizons in Figure 4.2 are required. As explained, the termination horizon is included to address the difficulty of computing terminal constraints associated with the terminal state of the building and its systems. To ensure that the initial conditions are the same for each new optimization, each simulation is run through an initialization period as described in [24]. The initialization horizon should not be confused with the warm-up period in EnergyPlus simulations,which precedes the initialization horizon in each simulation; rather, the initialization horizon ensures that long-term effects of prior control strategies are accounted for. We have found that a 7-day initialization horizon is sufficient to ensure that states are matched, even for buildings with high thermal mass.

A final note on horizons as they are treated in this application of SMPC is the way that control decisions are applied during the initialization horizon. For the first optimization at $t = t_0$, the initialization horizon for the default simulation and for all candidate simulations is the same, as indicated in Figure 4.3. Once the system advances in time however, the optimizer will inevitably choose different controls than what are used in the default case, and in subsequent optimizations, the control decisions applied during the initialization horizon diverge as in Figure 4.3. So, for all simulations after $t = t_0$, default cases use default control setpoints in the initialization horizon, and candidate cases use optimal setpoints in the initialization horizon; this ensures that any long-term planning by the optimizer can be taken advantage of, i.e. pre-cooling a large thermal mass on one day in preparation for a warm following day.

4.3.4 Cost Functions

In Equations 4.1 through 4.4, the SMPC problem is defined, and the objective function is defined, with attention given to how it is broken up into different components for time horizons. Here we describe how the objective function is broken up in terms of energy

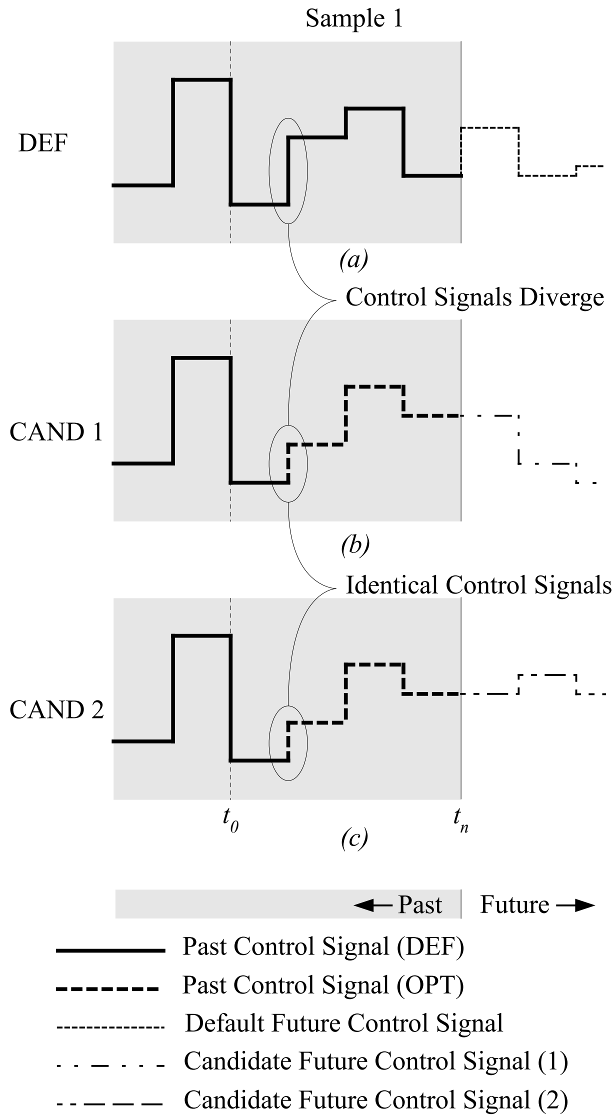

Figure 4.3: Divergent thermal histories: for optimizations at $t \geq t_1$, candidate simulations have the optimal thermal history, while default simulations have the default thermal history.

consumption and comfort penalties, and define the particular objective function, henceforth referred to as the cost function, used in the results presented later in this paper. The discussion begins with defining the primary cost function, used in the majority of simulation studies. At the end of this section, a secondary cost function is defined which is aimed at improving comfort conditions - rather than just respecting a discomfort limit.

In dealing with a suite of simulation results, computing the cost function is a two-step process. First, the cost function is computed for each simulation in the suite, then some method of aggregating the results must be applied to arrive at a singular numerical value that the optimizer can associate with each candidate control vector. The most appropriate statistic to use in this case is the expected value, or arithmetic mean of all of the individual simulation results. Another option for treating the range of results would be to take the maximum value, essentially assuming the worst and guaranteeing satisfactory albeit conservative controller performance.

Since the optimizer must be presented with a single numerical value, we use the expected value $E(C(\vec{u}))$ of building performance given each control vector $\vec{u}$, where the cost function $P(\vec{u})$ (or similarly, $T(\vec{u_T})$) is a metric for defining building performance for a given period of time, defined in Equation 4.5,

$$P(\vec{u}) = \sum_{t=t_0}^{N} R(\vec{u}(t)) + M(\vec{u}(t)) \tag{4.5}$$

where N includes all time steps in the planning horizon, R is the cost associated with energy consumption, and M is the cost associated with comfort performance.

While the cost of energy consumption is straightforward for computation, there is significant ambiguity around how to quantify comfort - especially when considering temporal datasets. Here the cost of comfort is defined using hourly PMV scores, occupancy, and a base-case for comparison, as presented in [24, 63], and below in Equation 4.6. In essence, this sets a prescribed range of comfort conditions that are allowed, which is whenever the hourly PMV is between ± 0.5, or closer to zero than what was achieved under default control, and results in penalty for any comfort conditions that fall out of the range.

$$M(\vec{u}(t)) = [(|PMV_{t,CAND}| > 0.5) - (|PMV_{t,DEF}| > 0.5)]\, N_{occ_t} W_M \tag{4.6}$$

That is, for each candidate control vector $\vec{u}$, a simulation with default control setpoints is conducted, and a simulation with the candidate control setpoints is conducted. The PMV score for each hour from the candidate results is compared to the default results, and whenever the candidate PMV score is worse than the default score, or worse than a predetermined baseline value of $(PMV \pm 0.5)$, it is multiplied by the number of occupants present and a large weighting coefficient W_M, which serves to enforce constraints on comfort conditions in simulation. For the results presented later, W_M was set to a value of 10^6, which causes the cost function to be very large for any simulation with even slightly poor comfort conditions; essentially preventing the optimizer from choosing controls that lead to any comfort conditions which fall outside of the prescribed range.

4.3.4.1 Secondary Cost Function

From the start of simulation studies with the RSF building, it was obvious that saving energy would be difficult since the building was already well designed and operating efficiently. When initial simulation studies showed what we expected (minimal energy savings), a new cost function was defined that would incentivize the optimizer to try to save energy and **improve** comfort, rather than trying to save energy and merely meet the comfort constraints set by the default controller.

Recall that first cost function consists of two components, an energy cost (raw energy consumed) and a comfort penalty. The comfort penalty for a candidate simulation is zero when comfort conditions are the same as, or better than, the comfort conditions in a default simulation, as determined by the PMV score. The second cost function consists of four components, namely energy consumption, a comfort score, an energy penalty, and a comfort penalty. Each of these components is computed for an entire planning horizon, rather than per-hour. This new formulation is given below in equation 4.7.

$$P = \frac{1}{2}\left(\frac{R_c}{R_d} + \frac{M'_c}{M'_d}\right) + W_R\big(R_c - R_d\big)^+ + W_{M'}\big(M'_c - M'_d\big)^+ \tag{4.7}$$

Note that a new formulation for the cost of comfort for a given hour is redefined as M' to distinguish it from the cost of comfort in the first cost function, M. The new cost of comfort for a given hour is simply the occupant-weighted PMV score for that hour, as in equation 4.8.

$$M'(\vec{u}(t)) = N_{occ_t} * PMV_t \tag{4.8}$$

The weighting coefficients, W_R and $W_{M'}$ are both set to 10^6 to heavily penalize any increase in energy comfort or any degradation of comfort conditions.

Some hypothetical energy and comfort scores are given in Table 4.2 below, to illustrate the difference between the first cost function and the second, just described. The difference between the two cost functions is subtle, but evident in the second and fifth rows of the table, where the second objective function shows a benefit from improving comfort conditions, and the first shows none.

Table 4.2: Example calculations of cost function values for a single cost horizon, assuming the default simulation has resulted in 1 unit of energy consumption, and a 1 unit comfort violation. Note that a comfort value closer to zero is better.

Default		Candidate		Cost	Cost
Energy	Comfort	Energy	Comfort	Function 1	Function 2
1	1	1.0	1.0	1.0	1.00
1	1	1.0	0.9	**1.0**	**0.95**
1	1	0.9	1.0	0.9	0.95
1	1	0.9	0.9	0.9	0.90
1	1	0.9	0.8	**0.9**	**0.85**
1	1	1.0	1.1	100001.0	100001.05
1	1	1.1	1.0	1.1	100001.05
1	1	1.1	1.1	100001.1	200001.10
1	1	1.1	0.9	1.1	100001.00
1	1	0.9	1.1	100000.9	100001.00

Table 4.2 also illustrates the sensitivity of both cost functions to poor comfort performance; if the candidate control results in even slightly worse comfort conditions, the cost function value will be very high, due to the high weighting coefficients. For the second cost function, this is also true for energy consumption, and the optimizer is quickly forced away from any control decisions that result in higher energy consumption.

Two nuances of the second cost function are a sensitivity to zero values, and an ability to trade between hours. Firstly, a value of zero for either energy or comfort score for the default simulation (which can happen on weekends when no occupants are present) will result in an infinite value of one of the first terms ($\frac{R_c}{R_d}$ or $\frac{M'_c}{M'_d}$), so some logic is included in the code to address this issue. Secondly, since each individual part of this cost function (R_c, R_d, M'_c, and M'_d) is computed for the entire planning horizon prior to combining into the final score (P), the cost function allows for hourly energy costs to be traded with each other, or for hourly comfort scores to be traded. The two penalty terms ensure that comfort and energy are not traded with one another, as evidenced by the high values in the last two rows of Table 4.2. The ability to trade comfort scores relies on the assumption that one occupant's discomfort present for five hours when the PMV score is PMV_{-1} is equivalent to five occupants' discomfort present for one PMV_{-1} hour.

4.3.5 Handling Stochastic Occupant Behavior Models

This section addresses the treatment of stochastic phenomena in occupant behavior within building energy models used here. Within the building energy model, a stochastic model of occupant window use is implemented which requires a random number at each time step to generate unique evolutions of occupant behavior. The seeding of this randomly generated number is an important detail since it governs stochastic phenomena within the model, and its importance is illustrated in Figure 4.4, where the different combinations of automatic controls and stochastic phenomena in simulation cases are illustrated.

For a single MPC optimization (for one execution horizon), the optimization will con-

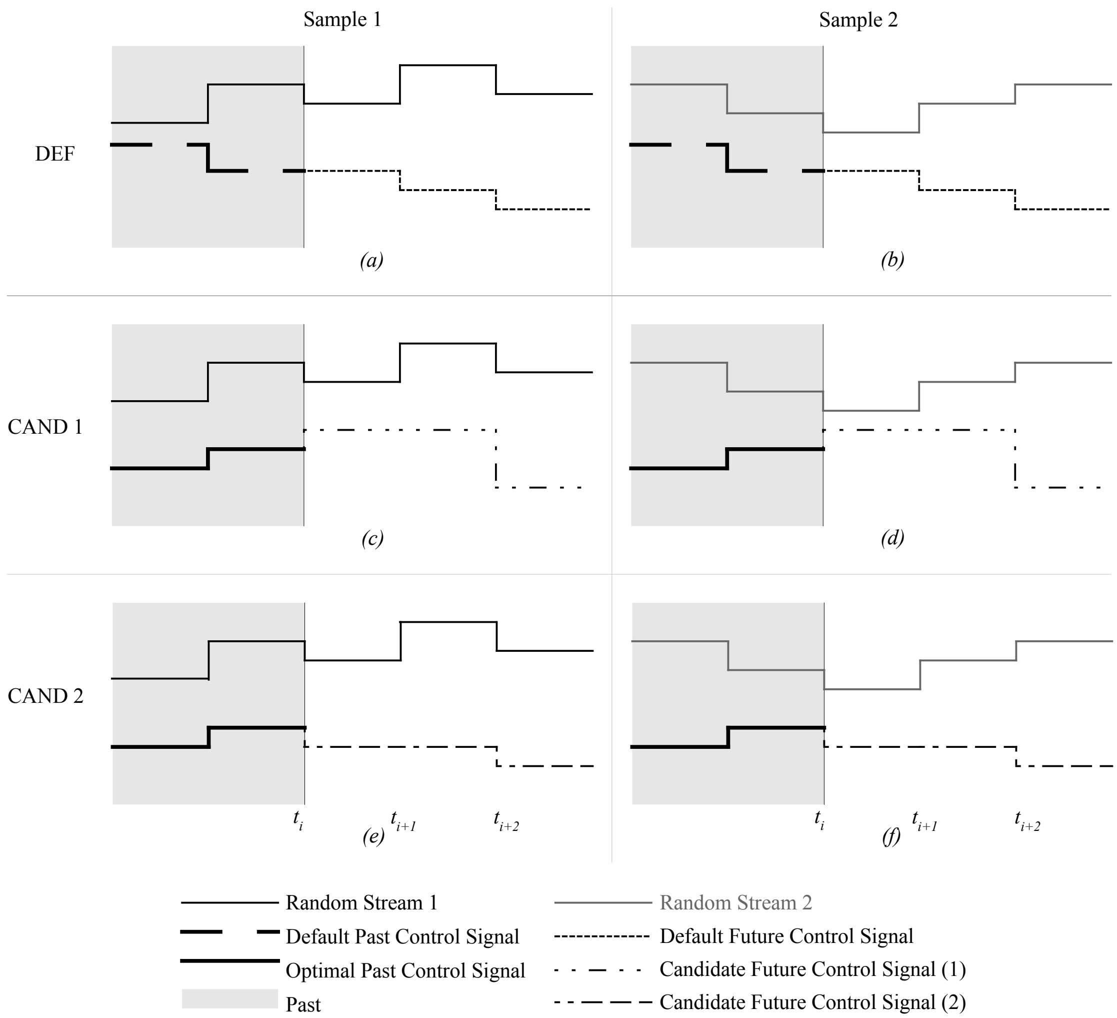

Figure 4.4: Illustrations of control signals and stochastic phenomena in different simulation cases.

duct a suite of simulations for the model with default controls (DEF), and suites of simulations for each candidate control vector (CAND) evaluated during the optimization, computing the objective function for each individual simulation, then aggregating the results of each suite.

In order to compare the results of a simulation for a given candidate control vector to the results of a simulation with default controls, the random seed must be the same in both simulations. This is shown in Figure 4.4 by the two solid thin lines labeled Random Stream 1 and 2, respectively. In the first column of subfigures (a,c, and e), the same random seed (Random Stream 1) is used in each simulation, and in the second column of subfigures, a different random seed (Random Stream 2) is used. When comparing different simulations, one can compare simulations from different rows or different columns of Figure 4.4, but not from different rows **and** columns. Comparing the simulations in Figure 4.4 (a) and (b) gives insight into the effect of different evolutions of stochastic behavior, given the same automatic (DEF) controls, while comparing the simulations in Figure 4.4 (a) and (c) gives insight into the effect of different automatic controls (DEF vs. CAND 1), considering a single evolution of stochastic behavior. It is not valid to compare the simulations in Figure 4.4 (a) and (d), since both the automatic controls **and** the stochastic behavior are different.

Note that this does not guarantee that the stochastic behavior will evolve in the same manner in each simulation with the same random seed, but that if all other parameters and the random seed are the same, then stochastic behavior will be identical in each simulation. In fact, one of the goals of this research is to find control strategies that work in concert with occupant behavior, so if a candidate control vector leads to indoor conditions which, in turn, lead to more desirable occupant behavior, it will be reflected in the results. If, on the other hand, a given control sequence leads to undesirable occupant behavior, this will be reflected in the results and the control sequence will be discarded.

4.3.6 Monte Carlo Criteria

During a given optimization, some number F of candidate vectors are explored, and for each candidate vector, a suite of simulations are conducted to establish a representative result. The quantity of simulations per-candidate can be established a priori as N_{MC}, as show in Table 4.3, or can vary with each candidate depending upon the variability in the stochastic behavior for that candidate, as in Table 4.4. Choosing a set number of Monte Carlo simulations a priori is difficult for two reasons, first it is impossible to know in advance how many simulations will be required to capture the range of variability in the stochastic model being simulated - so the chosen number may be too low; the second is that simulations are computationally expensive, and for cases with a small range of variability the chosen number may be too high. The solution to this conundrum is to implement a system of checking the results after any number of simulations have been conducted to determine if enough simulations have been conducted to arrive at a satisfactory expected value.

4.3.6.1 Convergence Band Stopping Criteria

Table 4.3: Monte Carlo sampling and optimization candidates - static Monte Carlo criteria.

	Sample 1	Sample 2	$\cdots$	Sample N_{MC}
Default	Def, S_1	Def, S_2	$\cdots$	$Def, S_{N_{MC}}$
Candidate 1	C_1, S_1	C_1, S_2	$\cdots$	$C_1, S_{N_{MC}}$
Candidate 2	C_2, S_1	C_2, S_2	$\cdots$	$C_2, S_{N_{MC}}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Candidate F	C_F, S_1	C_F, S_2	$\cdots$	$C_F, S_{N_{MC}}$

A variant of the methodology in [6] is employed to determine when sufficient simulations have been performed to ensure that the range of probable impacts of occupant action has been explored. First, a pair of convergence criteria: the convergence band width B_W and the convergence band length B_L are defined, as in Equation 4.9 and Figure 4.5. The band length is expressed as a minimum number of simulations, and the band width is a percentage

Table 4.4: Monte Carlo sampling and optimization candidates - variable Monte Carlo criteria.

	Sample 1	Sample 2	$\cdots$	$\cdots$	$\cdots$	Sample N_{MC}
Default	Def, S_1	Def, S_2	$\cdots$	$\cdots$	$\cdots$	$Def, S_{N_{MC,D}}$
Candidate 1	C_1, S_1	C_1, S_2	$\cdots$	$\cdots$	$\cdots$	$C_1, S_{N_{MC,1}}$
Candidate 2	C_2, S_1	C_2, S_2	$\cdots$	$C_2, S_{N_{MC,2}}$	$\times$	$\times$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Candidate F	C_F, S_1	C_F, S_2	$\cdots$	$\cdots$	$C_F, S_{N_{MC,F}}$	$\times$

error, E, of the average value of all simulations up to simulation n, analogous to a statistical confidence interval.

$$B_W = E\bar{C}_n \tag{4.9}$$

$$\bar{C}_n = \frac{1}{n}\sum_{i=1}^{n} C_i \tag{4.10}$$

After computing the cost function of simulation n, one can compute the upper and lower bounds of the convergence band as $\bar{C}_n \pm B_W/2$.

Next, we ensure that the moving average at each prior iteration falls within the convergence band, as shown by the dark line in Figure 4.5, which is within the darkest shaded region for all simulations in B_L. This criteria is also given in Equation 4.11. Values of B_L and B_W are 50 and .02, respectively, which ensures that a minimum of 50 simulations are conducted, and that the moving average does not change by more than 2% over the last fifty computations.

$$\left[\bar{C}_n - \frac{B_W}{2} \overset{!}{<} \bar{C}_{n-k} \overset{!}{<} \bar{C}_n + \frac{B_W}{2}\right] \forall k, k = 1, 2, ..., B_L \tag{4.11}$$

This ensures that when the impact of occupant behavior is small, superfluous simulations are not needlessly conducted, and when occupant behavior leads to large variation in the results, we perform enough simulations to capture that variability.

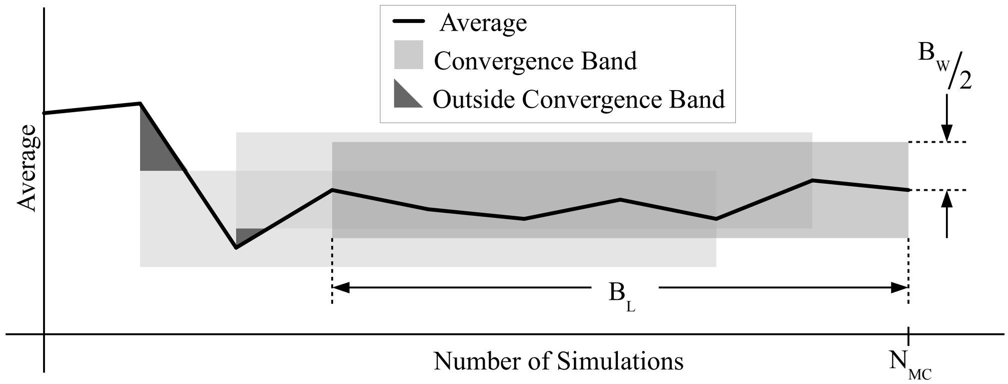

Figure 4.5: Convergence band method for Monte Carlo stopping criteria.

4.3.6.2 Variance Based Stopping Criteria

Following the methodology in Section 4.3.6, we begin with an investigation of normally distributed data, and demonstrate through empirical studies how we can ensure that enough simulations have been conducted in a Monte Carlo sampling exercise with simulated building energy consumption data. First, we recall the standard normal distribution $\Phi(z)$, with mean μ and standard deviation σ. According to the central limit theorem, the distribution of any random variable, x, will approach the normal distribution as the sample size, n, approaches infinity.

$$\lim_{n\to\infty} Pr\left(\frac{\bar{x}_n - \mu}{\sigma/\sqrt{n}} < z\right) = \Phi(z) \tag{4.12}$$

Using the standard normal distribution, one can compute the probability that the sample mean is equal to the population mean (in this case, equal to the mean of the standard normal distribution) according to

$$Pr\left(|\bar{x}_n - \mu| < z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\right) = 1 - \alpha \tag{4.13}$$

where $(1-\alpha)$ is the desired level of confidence, and $z_{\alpha/2}$ is equal to the inverse normal distribution calculated at the upper $100(1-\alpha/2)$th percentile as in

$$z_{\alpha/2} = \Phi^{-1}(1-\alpha/2) \tag{4.14}$$

By expressing the difference between the sample mean and the true mean as a relative error $\epsilon = E * \bar{x}$, we can determine the the number of Monte Carlo n_{CLT} samples required to say with $100(1-\alpha)\%$ confidence that the sample mean is within $100(\epsilon)\%$ of the true mean.

$$|\bar{x}_n - \mu| = \epsilon \tag{4.15}$$

$$\epsilon < z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \tag{4.16}$$

$$n_{CLT} = \left\lceil \left(z_{\alpha/2}\frac{\sigma}{\epsilon} \right)^2 \right\rceil \tag{4.17}$$

Note that we now have three parameters, σ, α, and ϵ which determine the number of samples required. Since σ is fixed at a value of 1.0 for the standard normal distribution, we can observe the number of samples required as a function of α and ϵ in Figure 4.6.

Figure 4.6 shows us that the number of samples required increases sharply as the error criteria shrinks. If we look at a normal distribution with mean $\bar{x} = 1$ and standard deviation S_x, and if we fix the error criteria to 0.1, we can look at the number of samples as a function of the standard deviation and the confidence level in Figure 4.7.

Note in Figure 4.7 that the number of samples goes up steeply as the standard deviation increases relative to the mean of the data, and that even for a modest 10% error ($\epsilon = 0.1$) and 95% confidence ($\alpha = 0.05$) levels, some 400 samples are required if the standard deviation is large relative to the mean ($S_x/\bar{x} = 1$). The sample size n_{CLT} is only applicable to independent and identically distributed (i.i.d.) random variables that come from a normal distribution, and ensures that the variance of the sample-data is very close to the variance of the normal distribution.

Next, we apply the standard sample size on an empirical data set, and show that it is overly conservative for our needs since it depends on the variance of the data, while we are only concerned with finding a representative mean.

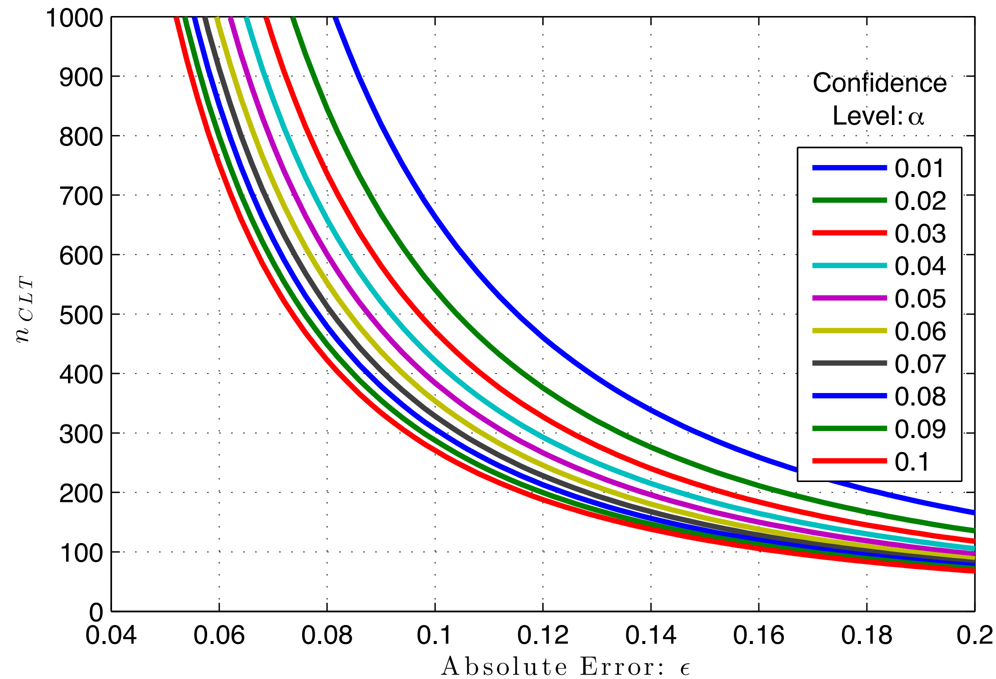

Figure 4.6: Monte Carlo samples required to reach representative mean for the standard normal distribution.

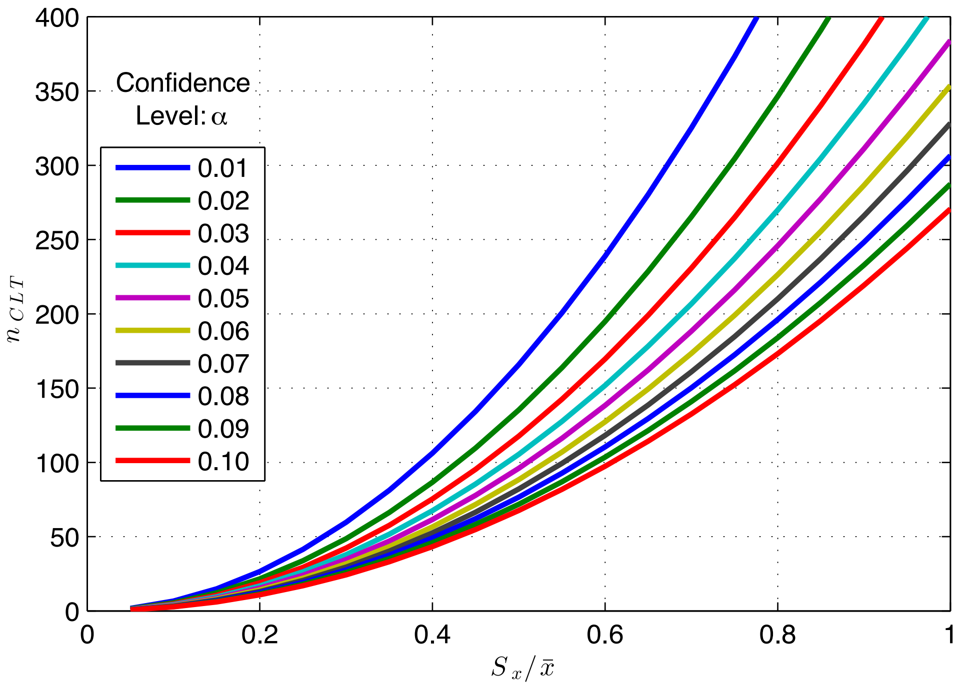

Figure 4.7: Monte Carlo samples required to reach representative mean for normally distributed data with mean $\bar{x} = 1$ and error criteria $\epsilon = 0.1\bar{x}$.

4.3.6.3 Empirical Testing

For empirical testing, 1000 simulations of the SMM1 model (described in Chapter 5) with stochastic occupant window behavior were conducted for the month of June, since it is conducive to occupant window use, and represents a case where occupant behavior heavily impacts building performance. Figure 4.8 shows the minimum and maximum energy consumption for each hour of the simulation.

In Figure 4.9, we see the distributions of daily variation in energy consumption for the month of June. Note that early on in the month, outdoor conditions are cooler and occupant use of windows is unlikely. Later in the month however, occupant window use leads to a wide range of energy consumption on some days. When we conduct SMPC, we consider the average energy consumption over the course of 8-10 days, depending on the horizon scheme implemented, which includes a 7-day initialization and 1-3 day cost horizon, so it is important that we capture the variability over a 10-day period. In the beginning of June, this variation is small - so theoretically we do not need to conduct very many simulations to get a good representative average of energy consumption. Later in the month however, occupant behavior leads to a wide spread in the data, and we will want to conduct more simulations to capture that spread.

Table 4.5 contains the statistics of the daily energy consumption data presented in Figure 4.9, with the mean and standard deviation from the full 1000 runs presented in the first two columns, and the required number of simulations needed to arrive at a representative average according to two methods. **CB Stop** refers to the number of simulations that are required to reach convergence using the convergence-band method described in Section 4.3.6.1, and **Variance Stop** is equivalent to n_{CLT} in Figures 4.6 and 4.7; it is the number of simulations required to reach a representative average using the variance-based standard sample size methodology in Section 4.3.6.2.

An example demonstrating the two convergence criteria for June 26 (CB Stop and

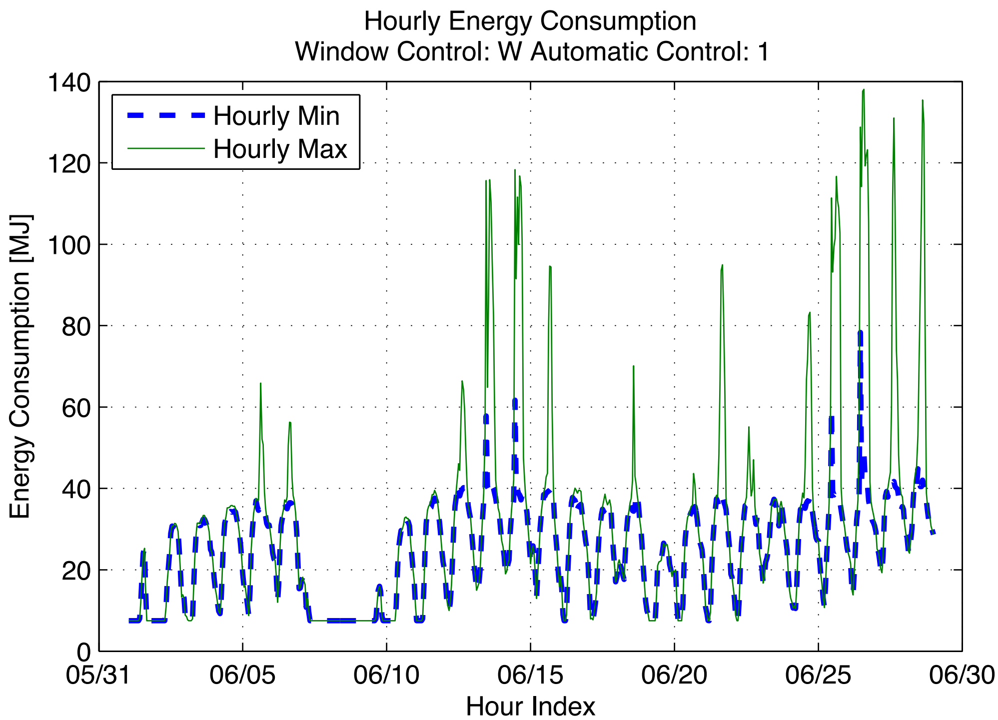

Figure 4.8: Hourly variations in energy consumption considering occupant use of windows for 1000 simulations with the SMM1 model for the month of June.

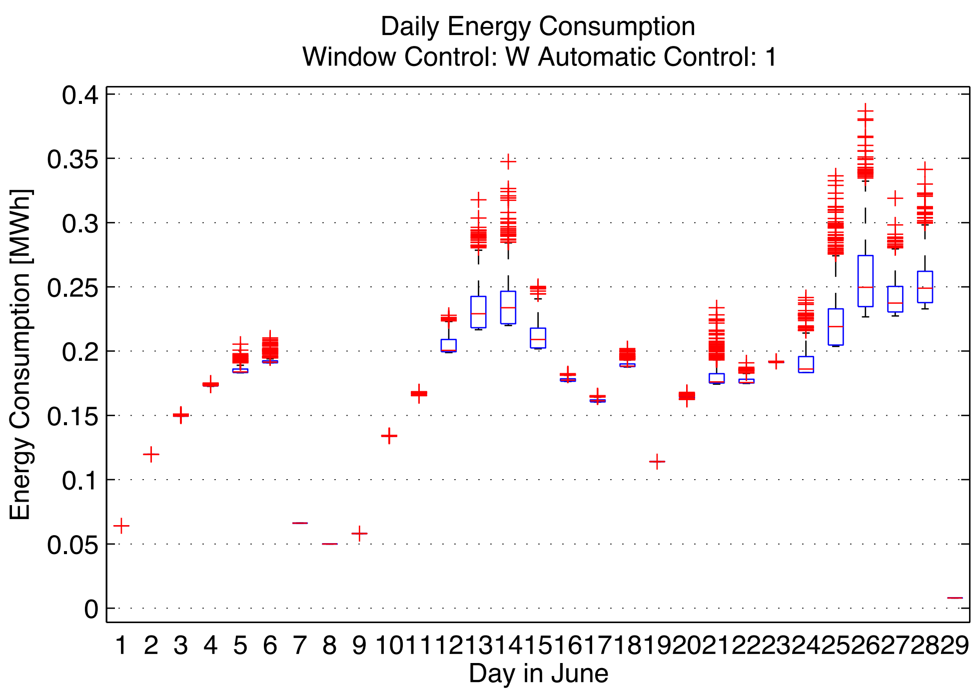

Figure 4.9: Daily energy consumption.

Table 4.5: Daily energy consumption statistics (kWh) and convergence criteria.

Day	Mean	Standard Deviation	CB Stop	Variance Stop
1	64	7.37E-06	36	1
2	120	0.015	36	1
3	150	0.168	36	1
4	173	0.338	36	1
5	185	3.145	36	12
6	193	3.239	36	11
7	66	1.58E-04	36	1
8	50	5.00E-13	36	1
9	58	3.72E-05	36	1
10	134	0.052	36	1
11	167	0.412	36	1
12	204	5.789	37	31
13	234	17.793	51	223
14	238	20.039	52	273
15	211	10.079	42	88
16	178	1.261	36	2
17	161	1.016	36	2
18	189	2.179	39	6
19	114	1.17E-03	36	1
20	164	0.620	36	1
21	180	8.610	42	88
22	177	2.621	36	9
23	191	0.213	36	1
24	192	10.956	71	126
25	225	24.556	71	458
26	258	31.009	71	554
27	241	14.276	51	135
28	253	17.949	78	194
29	8	1.63E-13	36	1

Variance Stop) is given in Figure 4.10. The confidence band and width parameters are 36 and 0.01 for the CB method; the absolute error and confidence level α are 0.01 and 0.05 for the variance method. Note that the according to the CB method, we have reached a representative average after 71 simulations, while the Variance based method ensures that both a representative average **and** a stable variance have been reached at 554 simulations. For the purpose of our SMPC optimizations, the CB method is sufficient, since it provides a representative average with the fewest simulations possible.

Table 4.6 has statistics for energy consumption on June 26 after 71, 554, and 1000 simulations. Assuming that the 1000 simulations encompass the entire population, and that the mean, standard deviation, and variance are their true values after 1000 simulations, we have computed the difference between the true and sample statistics after 71 and 554 simulations. Note that after 71 (the stopping point according to the CB method), the mean is within 1% of the true value, but the standard deviation and variance are still 5 and 11 % away from their true values. Following the variance based stopping method, one ensures that the standard deviation and variance are much closer to their true values, but the mean value is certainly captured by the CB method.

Table 4.6: June 26 energy consumption statistics [kWh].

Simulation Number	Cum. Mean	Cum. Std. Dev.	Variance	% Diff. Mean	% Diff. Std. Dev.	% Diff. Variance
71	256	33	3852	-0.86	5.36	11.28
554	260	31	3547	0.54	0.89	2.46
1000	258	31	3462	-	-	-

Next, we consider a ten-day simulation result, since this is what will commonly be visible to the optimizer. Since we are considering now 240 hours instead of 24, the impact of occupant actions in individual hours are much less significant, and the results that were much more variable on a 1-day time scale are washed out in the 10-day summary, presented

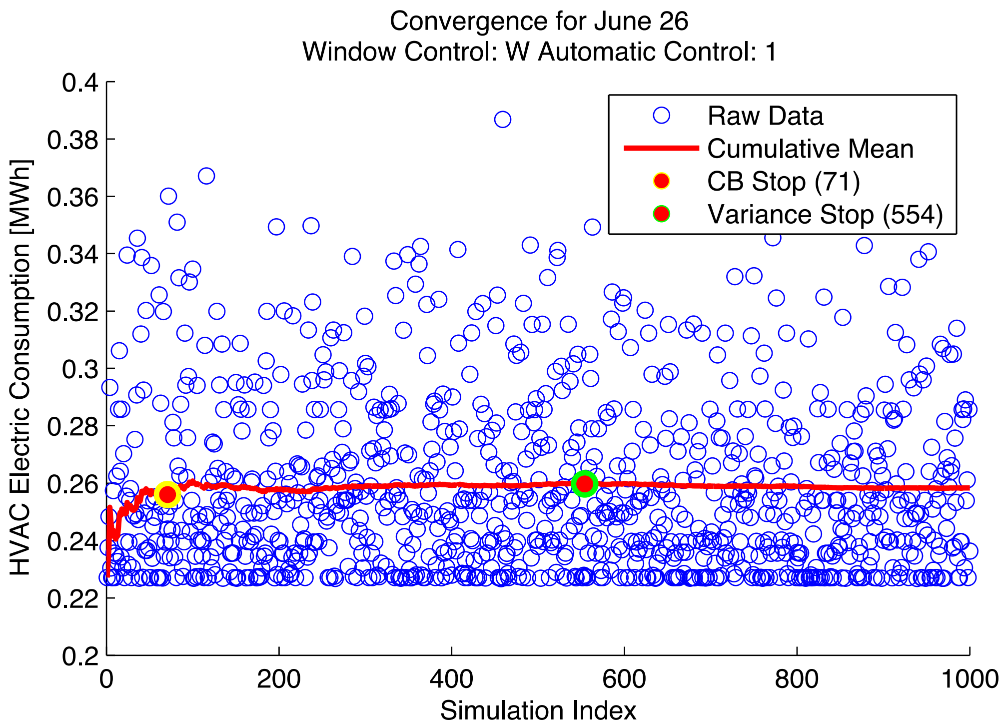

Figure 4.10: Individual simulation results, cumulative mean, and stopping criteria for June 26.

in Figure 4.11. Where the range of energy consumption impacts for a single hour was 300% of the minimum value, and the range for a single day was 50%, the largest distance between the minimum and maximum for a ten day period is only about 20%. This smaller range in values means that the rate of convergence to a representative average will be even faster for longer time-scales than for shorter ones.

If we look at the summary statistics for the first twenty 10-day periods in June in Table 4.7, it is clear that using the variance-based convergence criteria leads to convergence in fewer simulations than were required for single-day metrics. For hourly data, the CB method consistently required fewer simulations to reach convergence than the variance method, but for ten-day data, the CB method tends to require more simulations - until the last two ten day periods. The ten day periods starting on June 19 and 20 have the highest variance, and the CB method and variance-based method require comparable numbers of simulations to reach an acceptable mean value.

If we consider the entire month of data together and compute the number of simulations required to reach an acceptable average according to each of our two rules in Table 4.8, we see that the variance based method requires only 16 simulations, and the CB-method requires only two more than its minimum, or 38 total.

In conclusion, the convergence band and variance-based methods give comparable stopping criteria when tested on a large (1000-sample) dataset for ten-day periods, which are typical time-periods in the MPC scheme described earlier in this chapter. While the variance based stopping rule correctly recommends stopping after just one or two simulations in many cases, it requires knowledge of the variance of the entire population of data in order to confirm that the single result is representative of the population. In contrast, the convergence band stopping rule does not require statistical knowledge of the population in order to establish convergence, but it does require a minimum number of samples before converging. Since we can not compute 1000 simulations every time a new control strategy is tested, and we need a representative average result with the minimum simulations possible, we select

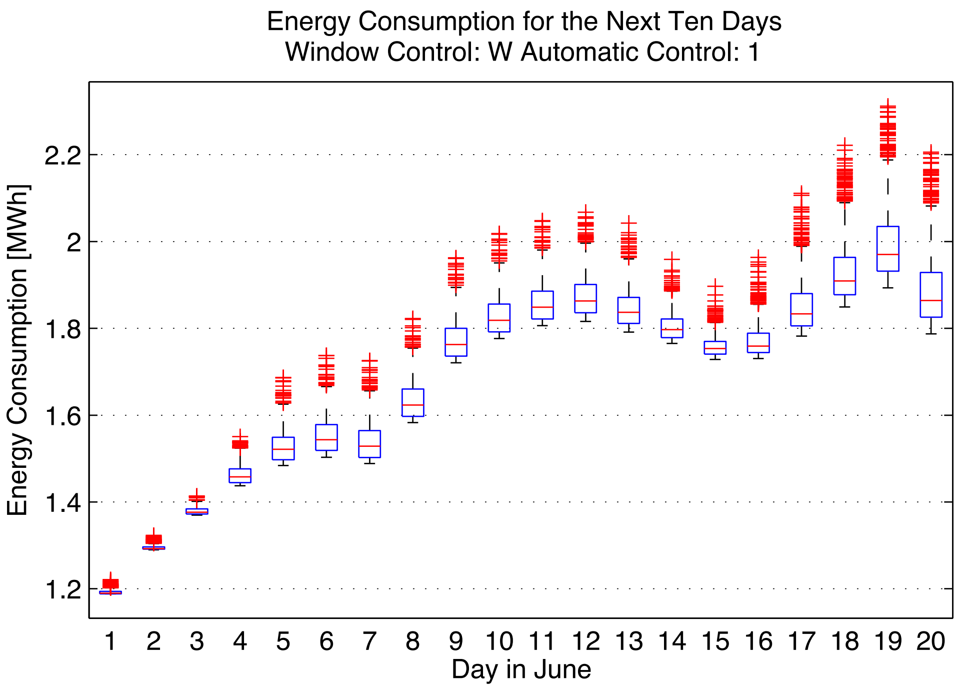

Figure 4.11: Variations in energy consumption for 10-day periods in June.

Table 4.7: Ten day energy consumption statistics (MWh) and convergence criteria.

Days	Mean	Standard Deviation	CB Stop	Variance Stop
1 - 10	1.193	0.0056	36	1
2 - 11	1.295	0.0056	36	1
3 - 12	1.379	0.0080	36	2
4 - 13	1.464	0.0222	37	9
5 - 14	1.528	0.0375	43	24
6 - 15	1.554	0.0436	43	31
7 - 16	1.539	0.0442	43	32
8 - 17	1.635	0.0447	43	29
9 - 18	1.774	0.0455	43	26
10 - 19	1.830	0.0455	43	24
11 - 20	1.860	0.0455	43	24
12 - 21	1.874	0.0468	43	24
13 - 22	1.847	0.0444	43	23
14 - 23	1.804	0.0319	39	13
15 - 24	1.758	0.0237	37	8
16 - 25	1.772	0.0385	42	19
17 - 26	1.853	0.0626	53	44
18 - 27	1.933	0.0739	51	57
19 - 28	1.996	0.0865	70	73
20 - 29	1.890	0.0865	70	81

Table 4.8: Monthly energy consumption statistics (MWh) and convergence criteria

Mean	Standard Deviation	CB Stop	Variance Stop
4.78	0.097	38	16

the convergence band method for use in future studies.

4.3.7 MPC Software Environment

The core of the simulation/optimization environment that is used to orchestrate simulated SMPC is described in detail by Corbin et al. in [24]. A summary is included here for clarity, see Figure 4.12 for a diagram of how data flows between software during a simulated MPC run; the numbered ovals indicate the sequence of events in the diagram. In Figure 4.12, an MPC run begins with the MPC Controller, which is implemented in MATLAB, which calls for an occupancy schedule from the occupancy model (discussed on Section 5.3.2). This occupancy schedule, which contains hourly presence and absence profiles for each occupant in the simulated model, will be used for the duration of the study. Next, the MPC controller gathers the data required to conduct a single optimization for a single execution horizon, and dispatches the information to an optimizer, (see number 3 in Figure 4.12). When the optimizer has concluded its work for a given execution horizon, the MPC controller stores the results of simulating with the optimal control vector u^* and advances to the next time horizon, repeating this process (all of steps 2 through 7 in Figure 4.12) until a predefined end-date is reached. Note that there are three nested processes in the diagram; the iterative simulation between steps 4 and 5 occurs in each objective function call, there are many objective function calls (steps 3-6) in each optimization, and one optimization is conducted for each execution horizon (steps 2-7) - resulting in a computationally expensive process.

4.3.7.1 Extending the Environment

While the focus of this work is largely focused on optimizing natural ventilation by scheduling automatic windows to open and close at certain times, the software environment is extensible to any parameter or setting in EnergyPlus. Since EnergyPlus uses text-based input data files (IDFs), and the software environment overwrites selected bits of text in the IDF in order to manipulate the simulated building controls, any field that is an input in the

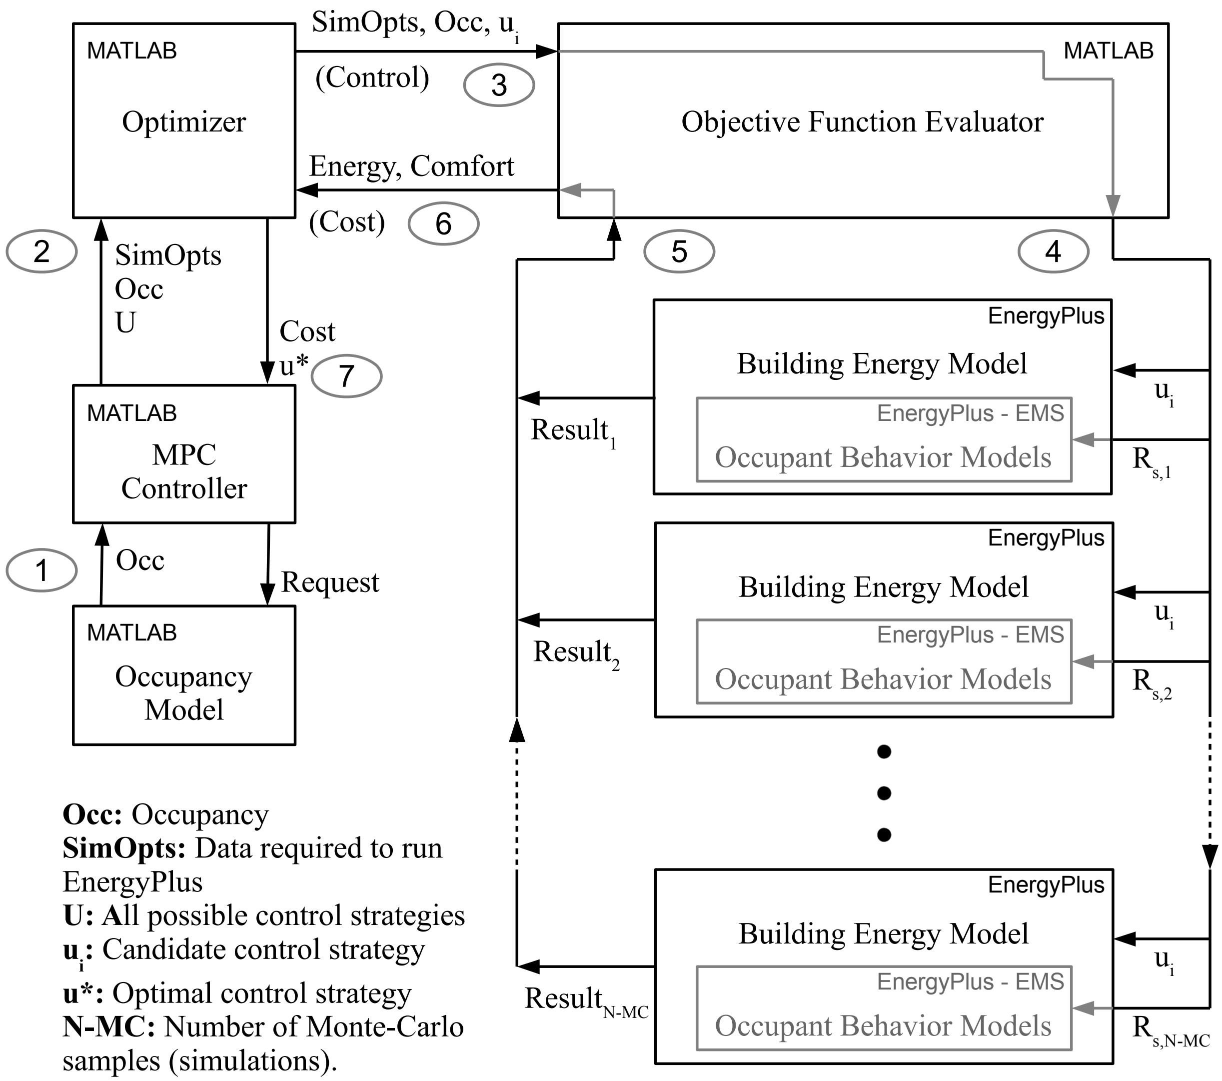

Figure 4.12: Offline SMPC software architecture.

IDF could be a parameter tuned by the optimizer. That said, we could in theory optimize schedules for lighting, HVAC setpoints, automatic window, or automatic blinds all at once in a single optimization study.

However, before embarking on an even more computationally expensive MPC study where every system is optimized together, consider that each system could have somehow conflicting objectives, and constructing a single objective function that will capture the goals of all systems at once will be a challenge. Consider also the granularity of parameters to be optimized in both scale and time - do we want to adjust lighting levels by 1 lumen at 10 second intervals, or by 50 lumens at 10 minute intervals?

Currently the software environment is limited to making hourly decisions, since any finer resolution would have increased the computational time beyond reasonable limits, assuming EnergyPlus is the simulation engine. The final optimization runs, discussed in Section 9.4, required 7 days of computing time using 12 2.8Ghz processors. Each run consisted of 213 individual one-day optimizations, each objective function call required between 36 and 100 simulations, and the only optimized parameters were 12 binary window open or close decisions. If more variables are to be considered for optimization, or continuous variables are considered rather than binary, the time for the optimizer to converge over each optimization horizon will increase exponentially.

4.4 Rule Extraction

In the top-down approach, the methods of machine learning adopted by May-Ostendorp et al. [63], are used to learn what relationships exist between ambient conditions and optimal control setpoints.

Here, the term **rule extraction** refers to the process of deriving usable building control rules from synthetic optimal datasets, such as those generated during an offline DMPC or SMPC process. In the context of this research, offline SMPC studies will yield large time-series datasets that include building performance characteristics, ambient conditions, and a set of optimal control decisions. Using various machine learning techniques, functional relationships can be developed between any subset of results (e.g. weather conditions) and the optimal control decision, which can then be converted into usable building control logic. Based on the work of May-Ostendorp et al, in which generalized linear models (GLMs), classification and regression trees (CARTs) and adaptive boosting methods were employed, we know that rule extraction works in an offline context.

4.4.1 Classification and Regression Trees

Rooted in the fields of data mining and decision theory, classification and regression tree analysis provides a nonparametric means to relate a set of predictor variables to a response variable through a sequence of conditional statements. CART models can be visually represented and understood through dendrograms (see Figure 4.13), and thereby lend themselves to implementation in building automation systems as simple if/then/else control rules. Regression trees are created by recursively splitting the set of predictor variables according to their skill at classifying the response variable; this process is referred to as learning or **growing** the decision tree. At each split (node) in the tree, the subset of predictor variables is examined, and a single predictor is selected that divides the subset into groups (branches) that similarly classify the response variable. When classification or misclassification rate is the sole metric used to determine where to split the predictor set, this process continues until every branch leads to a terminal node where there is one and only one data point, called a leaf. Normally, a fully grown tree (one in which every branch ends in a single value) is too complex and over-fitted to the data and must be simplified (pruned) to be useful. In order to prune the tree appropriately, a complexity criterion C_α is introduced that combines the rate of misclassification $R(T)$ with the number of terminal nodes (N_m) in the model. The misclassification error is given by:

$$R(T) = \frac{1}{N_m} \sum_{i \in R_m} I(y_i \neq k(m)) \tag{4.18}$$

as in [42], where I is an indicator function that returns a 1 when the expression in parenthesis is true, y_i is the the true class of a terminal node, $k(m)$ is the predicted class, and R_m is a region encompassing i observations that are isolated in a given split. The complexity criterion is given by:

$$C_\alpha(T) = \sum_{m=1}^{|T|} N_m Q_m(T) + \alpha|T| \tag{4.19}$$

where $T \in T_o$ is a subtree that is a obtained by pruning T_o, m is the index of a terminal node, α is a tuning parameter, and Q_m is an average of the squared difference between the total number of nodes and the number of nodes in the full tree T_o. The reader is referred to [42] for more information on the derivation of CART classification and pruning, suffice it to say that parsimonious models are readily found by established and effective pruning techniques.

4.4.1.1 Guiding Growth

The basic algorithm and premise behind CART growth is to minimize misclassification. For some data sets this is sufficient and the basic algorithm suffices, but the avid CART grower has several tools for guiding models in different directions. A maximum complexity value or a minimum number of observations per leaf can determine how complex a tree can grow, or one can specify that a certain number of observations must be present to even consider growing a new branch. Since misclassification is the main driver for tree growth, there are three tuning parameters that can be used to adjust the importance of misclassification for certain classes of observations or individual observations. Weightings are used to add importance to individual observations, while prior probabilities and class losses are used to add importance to classes of observations.

To illustrate, consider a dataset that has some multinomial response with three classes: 1, 2, and 3. A weight could be used to add importance to **individual** instances where the response takes on the value of 3, while a prior probability could add significance to **every** observation where the response takes the value of 3. Losses are more nuanced, because they can impart a conditionality to the importance of different observations. With losses, consider the true response, and the CART-predicted response; a perfect model would always predict a 2 to be a 2, a 3 to be a 3, and so on. Actual CARTs will misclassify some observations, and losses are used to penalize different misclassifications, so we could add significance to 2's that are misclassified as 1's, or take away significance from 2's that are misclassified as 3's.

In the end, using any of the three tuning methods can guide the CART to the same result. For a binomial response, specifying case weights of 2 for the first class and 3 for the second class is equivalent to specifying priors of 0.4 and 0.6, or specifying losses of misclassifying the first class as the second and misclassifying the second class as the first, as 1 and 1.5, respectively. Later in this research, priors are used to guide tree growth towards achieving a maximum ranked probability skill score (RPSS).

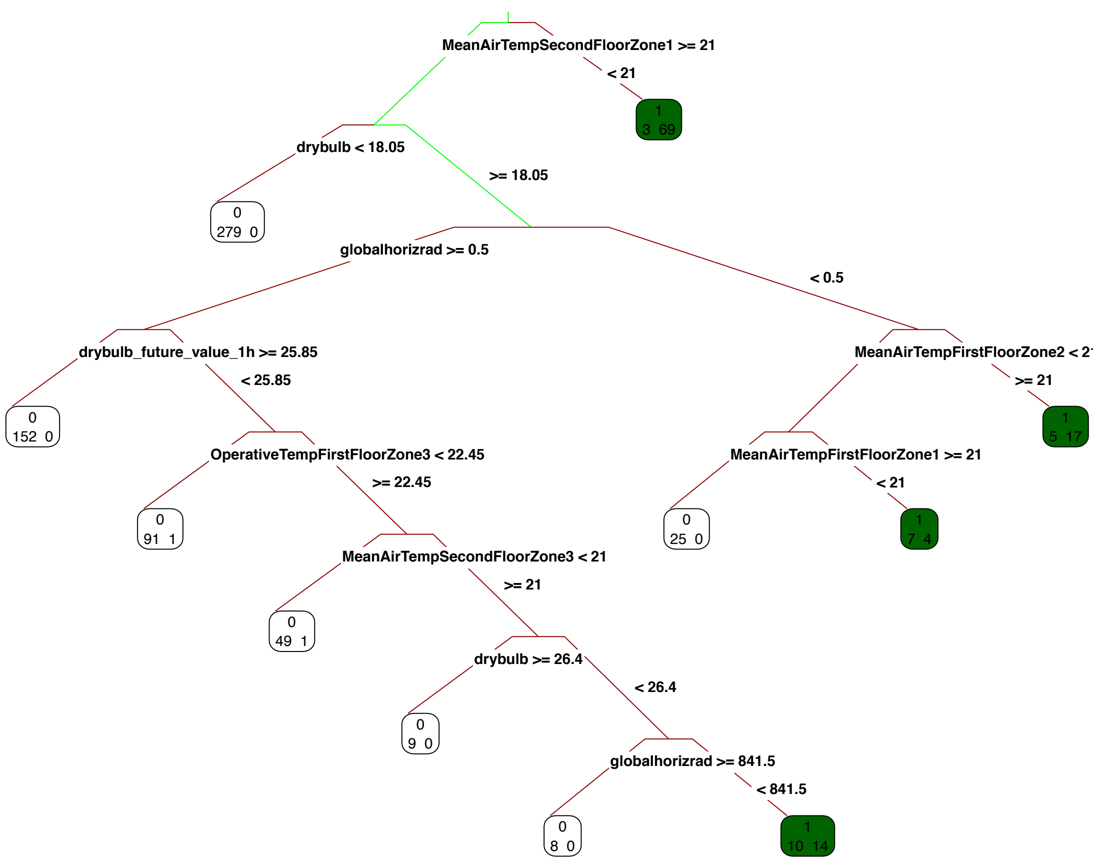

Figure 4.13: Full CART model dendrogram. The lighter green branches represent the subtree that would remain after applying the 1-SE pruning technique, eliminating the dark red branches.

Chapter 5

Model Development and Integration

5.1 Occupant Behavior Implementation

In general, occupant behavior models provide a probability at some instant in time of an occupant taking an action; the probability of an action P_A is computed as a function of some combination of current ambient conditions (C_{Env}) and occupant characteristics (Occ).

$$P_A = f(C_{Env}, Occ) \tag{5.1}$$

Given this probability and a randomly generated number (R_U) from the uniform distribution on the interval [0,1], the two numbers are compared to determine whether the action $A(t)$) takes place or not.

$$A(t) = \chi(P_A > R_U) \tag{5.2}$$

where $\chi(\cdot)$ is an indicator function that equals 1 if the statement in parenthesis is true. Choosing a threshold value instead of a random variable for comparison leads to a deterministic model of behavior, one in which every simulation would return identical sequences of behavior and thus energy consumption. When stochastic models are employed in simulations, they lead to uncertainty in simulation results, thus multiple simulations must be conducted to arrive at a representative distribution of results.

Whether it takes the form of a Markov chain, a bernoulli process, a logistic function, or something else, an occupant behavior model can be encoded as an algorithm. Since there is no standard means of representing occupant behavior in simulation or a standard means of modeling occupant behavior, different researchers have each chosen a unique language or approach to simulate their models. Here we describe two techniques that are used in this project to couple occupant behavior models to building energy models, specifically those built in EnergyPlus.

5.1.1 BCVTB

The Building Controls Virtual Test Bed [100], enables co-simulation between multiple simulation environments, including Modelica, Radiance, Simulink, and several others. As it pertains to this research, BCVTB enables occupant behavior models written in MATLAB® to interact with building energy models in EnergyPlus.

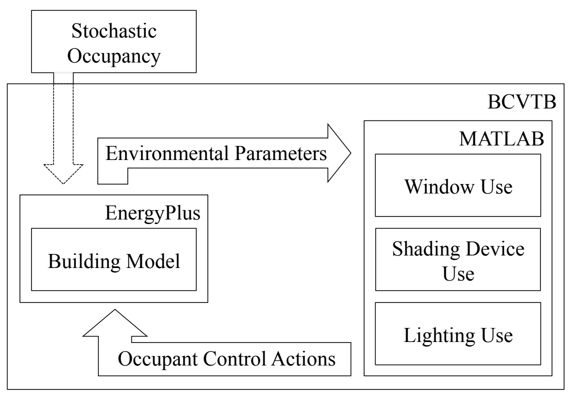

Figure 5.1: Building energy model and occupant behavior model connections.

5.1.2 EMS

In order to simulate a wide variety of controls and to enable users to customize models, EnergyPlus incorporates multiple objects in a subsystem called the Energy Management System (EMS). The EMS allows users to manipulate the state of virtually any simulation input, schedule, or control parameter. As of EnergyPlus version 7.0, the EMS also includes a built-in random number generator - which is a key ingredient to any stochastic model.

5.2 SMM1 Building Model

For initial investigations on the impact of occupant behavior in Chapter 6, an 11 zone EnergyPlus building model representative of a small mixed-mode office building is used. The roughly 1800 m^2 (18000 ft^2) model is pictured in Figure 5.2 below and includes manually operable windows, window shades, and lights. The MM model is derived from the DOE reference small office building model, and draws its base features and characteristics directly from the reference model; construction details are consistent with energy efficiency standard ASHRAE 90.1, the only building-controlled MM feature is building controlled windows. The building is oriented such that the longer exterior walls face North and South. The first level is divided into four perimeter zones and one core zone, while the upper two levels are divided into three zones: narrower zones on the east and west facades, and a large central core zone spanning the width of the building. For reference, this energy model is the same as the SMM1 model developed in [63].

Occupant density, lighting density, and plug loads are 0.0538 (persons/m2) 10.76 (W/m2), and 8.07 (W/m2) respectively. A conventional gas-fired heating system and direct expansion cooling systems serve the heating, ventilating, and air-conditionings needs of the zones through variable air volume terminal units when natural ventilation is not appropriate. Typical meteorological year data for Boulder, Colorado is used an all simulations.

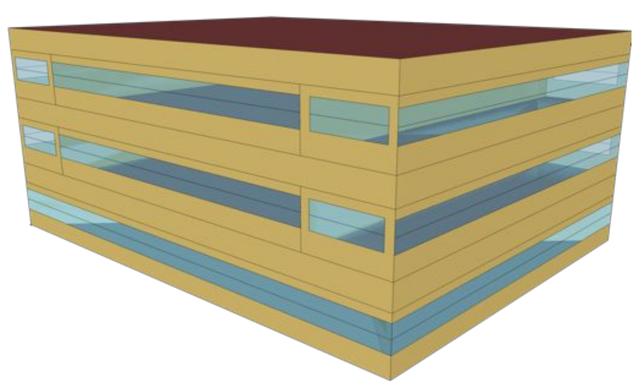

Figure 5.2: Isometric view of building energy model.

5.2.1 Incorporation of Occupant Behavior Models

In initial work, the BCVTB was used to couple the building energy model in EnergyPlus to occupant behavior models scripted in Matlab. When the quantity of Monte Carlo simulations grew however, and SMPC runs were beginning, it became evident that the addition of several extra input files and an additional simulation engine (BCVTB), were overly complex and significantly slowing down our simulation studies. At this juncture, the decision was made to script occupant behavior models directly into the building energy model via the EnergyPlus EMS.

Whether BCVTB or EMS was used, four occupant behavior models were used with the SMM1 building model: the Page model for occupancy, the Haldi models for window and shading device use, and the Lightswitch2002 model for occupant use of lighting only.

The Page occupancy algorithm generates an annual sequence of occupant presence, in its base form, the model predicts occupancy at 15-minute intervals for an entire year for a single occupant. Since the window, shading, and lighting use models require occupancy status as an input, and we assume there is one occupant responsible for each window, shade, or lighting device, the Page algorithm is used to generate individual occupancy profiles for each simulated occupant in the model. In the SMM1 model, there are only ten zones with external windows and shades, and a single occupant profile is simulated to 'control' each zone. For larger spaces where loads are coupled to dozens of occupants, the Page algorithm can be used to generate occupant presence/absence on a per-occupant bases, then aggregated to get a bulk percentage of occupancy at each time step.

To incorporate the occupancy schedule into EnergyPlus, the schedule has to be transformed in two ways; first, since we include the schedule as a separate input comma-separated value (CSV) file, it is limited to 8760 rows, or one per-hour, so instead of simulating presence at 15-minute resolution, it is simulated at hourly resolution. The second transformation is required for use with the window and shading device use models, which require not only occupant presence/absence - but also flags to indicate whether an occupant has just arrived, is about to depart, or has been present for a long time, and a few other parameters; there are 7 in total. These flags are pre-computed outside of EnergyPlus, and also included in the input schedule file.

In the SMM1 model, occupant behavior is simulated on a per-zone basis; that is, in spite of the fact that there would be dozens of occupants in each of the zones in the model, only one is assumed to activate the building systems, and we assume in all cases that a system is either fully activated (i.e. windows are fully opened) or fully deactivated (windows are closed). So for a single zone, when an occupant is predicted to open a window, all of the windows in that zone are assumed to fully open. This represents an extreme case, which is appropriate for investigating the range of impact that occupant behavior can have on building performance.

A last note on the implementation of occupant behavior models in SMM1 is that the average or default parameters for each behavior model are used. Often, models include tuning parameters to simulate more or less active occupants, and to account for diversity within the occupant pool. In this work, we only consider the 'average' occupant, and through a Monte Carlo analysis we learn what the average impact from the average occupant is.

5.3 NREL Research Support Facility

The Research Support Facility (RSF) (see Figure 5.4a) is the largest office building on the National Renewable Energy Laboratory (NREL) campus in Golden, CO. Even with a large data center on site, the RSF documented net-zero energy performance on June 23, 2011 with the help of clear skies and a solar array that covers the building and an adjacent parking structure. The RSF was designed to operate at 35kBtu/ft^2 (25kBtu/ft^2 without the data center) annually, and has met that achieved that level of performance. Ventilation in the RSF is provided by an underfloor air distribution (UFAD) system, controlled to maintain ambient CO_2 levels per ASHRAE 62.1-2010 [5], and by natural ventilation that is supplied

Figure 5.3: Research Support Facility (RSF) building in Golden, Colorado.

through both manually and automatically controlled windows, as illustrated in Figure 5.4b. Heating and cooling are provided by 42 miles of radiant piping embedded in concrete slabs, which comprise the exposed ceiling of each office space. With a gross area of 360,000 ft^2 and five floors, the RSF also features aggressive daylighting measures including external fins and shades, internal light-louvers, electrochromic glazing, and narrow wings (60 ft) with long north and south facing façades to minimize east and west facing glazing, maximize cross-ventilation potential, and to ensure that all interior spaces are within 30 ft of a window.

(a) RSF front entrance.[1]

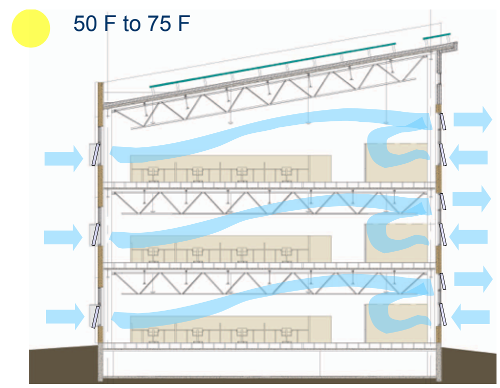

(b) RSF airflow concept drawing.[2]

Figure 5.4: RSF images.

5.3.1 RSF Building Model

For simplicity, a representative zone of the RSF building was modeled for use with control investigations; the window controls to be optimized only affect the open-plan office spaces of the building, so there is no need to model the data center, conference rooms, or other auxiliary spaces that are served by separate heating, ventilating, or air conditioning (HVAC) systems. EnergyPlus was selected as the modeling tool and simulation engine due to its ability to accurately model radiant systems and natural ventilation.

[1] http://www.nrel.gov/sustainable_nrel/rsf_photos.html
[2] http://apps1.eere.energy.gov/buildings/publications/pdfs/corporate/ns/webinar_rsf_03182010.pdf

The single zone model's physical dimensions, material constructions, schedules, loads, and systems are consistent with open-plan office spaces in the RSF, Figure 5.5 provides an image of the full building model, as well as the simplified 1-zone model used in this study. The RSF building was built using a modular design, in which 9.14m (30 ft) sections of the exterior walls were prefabricated, resulting in repeated identical building sections; the single-zone model represents two of these sections. The model also includes manually operable windows, automatically controlled windows, a concrete ceiling slab with embedded radiant piping for heating and cooling, an under-floor air distribution (UFAD) system for ventilation and supplemental conditioning, and an air flow network to account for natural ventilation. Heating and cooling for the RSF are provided by campus heating and chilled water systems, so purchased heating and cooling energy were specified in the model. The campus heating system burns wood waste for fuel and is 90% efficient, and the campus chiller plant has a COP of 7.8. Shading in the RSF is provided by fixed exterior shading elements, and internally mounted light reflectors that project daylight onto the ceiling. With the exception of the North and South facing (glazed) walls, all exterior surfaces of the single-zone model are assumed to have adiabatic boundary conditions, following from the assumption that this representative zone is surrounded by other similarly conditioned spaces. Schedules for lighting and electrical loads were adjusted to match published RSF-energy consumption data, schedules and controls for HVAC equipment were set to match existing building operation, and occupancy schedules were generated using the occupancy algorithm by [75]. Indoor conditions in the RSF are fairly constant due to the large thermal mass associated with the building's radiant systems, with indoor temperatures holding fairly constant at 23°C, which is well above the 10°C heating setpoint, and below the 24°C cooling setpoint.

5.3.1.1 RSF Model Calibration

Annual energy performance metrics (normalized by floor-area) as computed by the model and as measured on-site are presented in Figure 5.6. Note that the simplified model

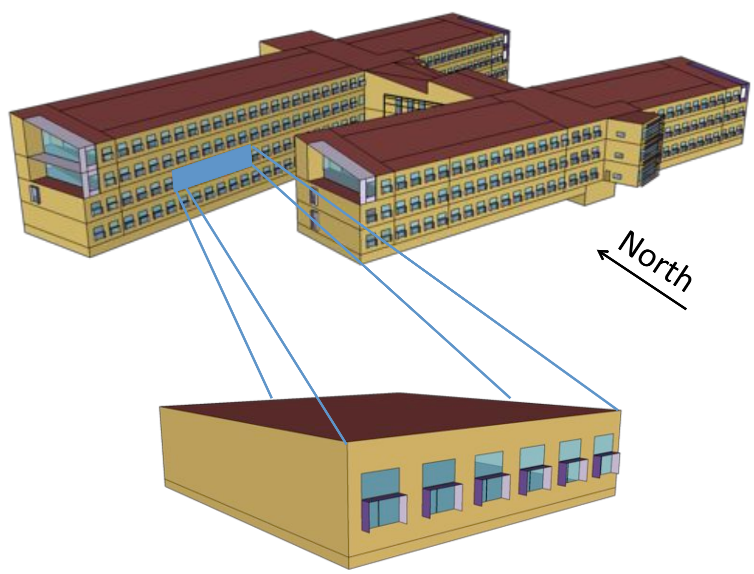

Figure 5.5: Simplified building energy model.

does not include loads associated with auxiliary spaces, and only predicts energy consumption for the open-plan office spaces, while measured data is for the entire building excluding the data-center. The model slightly under-predicts energy consumption in all areas; any savings predictions will be conservative.

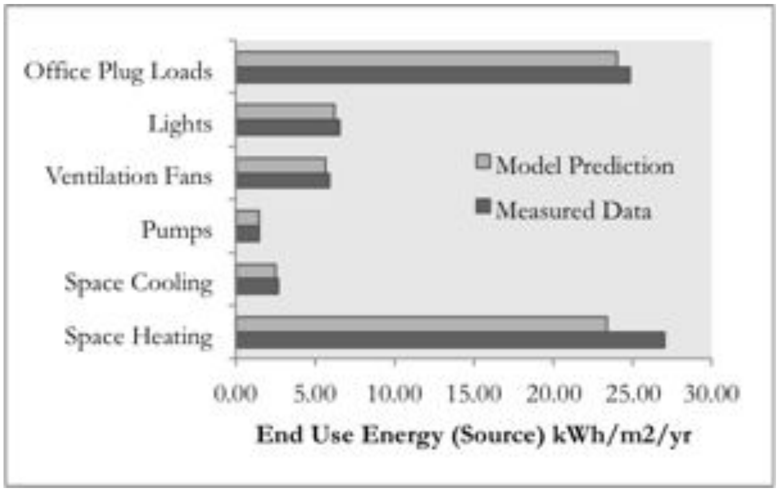

Figure 5.6: RSF measured energy consumption compared with simplified model predictions.

In developing and calibrating the EnergyPlus model, manual window operation was ignored, because the published data being matched was collected during the initial year of building occupancy, when the building's natural ventilation system was set to keep windows closed during occupied hours, and occupants were instructed to open and close windows consistently with the automatic windows. Average daily profiles of lighting and plug loads were available, so those loads are set as constant schedules in the model. Actual occupancy data from the Page algorithm is only used to inform the model of occupant window use during optimization studies.

After first developing the model to match parameters found in design documents, the primary steps taken to tune the model to match energy consumption data are outlined here.

- Operational schedules and setpoints were adjusted to keep indoor temperature fairly constant at 23°C year-round during occupied hours only.

- Window U-values were increased and wall U-values were decreased slightly from design values to raise the cooling loads and lower the heating loads in the model. Note that the assumption of adiabatic boundary conditions on four sides of the model means that the two external walls have a large impact on modeled performance, and in some sense serve as the envelope component for all sides of the model.

- Fan and pump sizing parameters were adjusted to appropriately serve the floor area of the model, and to consume the same power per unit area as the entire building.

A final note on the model is that detailed operational data was unavailable for the first two years of work on this project, so the model was tuned to match annual energy consumption. As more detailed data became available and as the actual control strategies for systems within the RSF building evolved, the model was changed to better match the best available data. The most significant changes to the model included changes to natural ventilation control logic as the physical building controls changed, and updating radiant system controls in the model to match the dynamics of the radiant systems in the physical building. The first iteration of the model included radiant system controls that allowed for fast changes in radiant slab temperatures, whereas radiant controls in the later iterations of the model are much slower.

5.3.2 Occupant Behavior Models

5.3.2.1 Manual Window Use

One model of occupant behavior applied in the case study here is described by Haldi et al. in [39], and determines whether occupant-controlled windows are open or closed for each time step in a simulation. The model is based on over seven years of occupant-window-use data from the LESO research facility in Lausanne, Switzerland; the reader is advised that the model may be unique to this climate, building, and set of occupants, however it remains the state of the art in the literature. At each time step during the simulation, the algorithm yields a probability P that an occupant will open or close a window, and if that probability is greater than a random number drawn from the uniform distribution on the interval $[0, 1]$, then the action occurs (in simulation). Once an action has occurred, the duration of that action persisting (i.e. how long a window stays open) is predicted by a sub-model which draws times from a Weibull distribution

The Haldi window model is implemented via the Energy Management System (EMS), and directly coded in the EnergyPlus Energy Runtime Language (ERL). In this implementation, the parameters corresponding to the 'average occupant' are used, these are the parameters derived from the aggregate dataset of window openings from all occupants in the behavior study. While there are parameters given for unique occupants in the behavior study, and the potential exists to model occupants with unique behavioral model parameters, that amount of detail was not included in this work. The reader is referred to [78] and [41] for detailed discussions on how occupant diversity impacts simulation results. To summarize, the model predicts a probability of window opening or closing based on indoor and outdoor temperature, the presence of rain, and on the state of occupancy, and the coefficient that corresponds to each ambient condition can vary from one occupant to the next. In the case study results presented here, the coefficients are constant for all occupants, as they were derived from an aggregate dataset for all occupants.

For example, with no rain, an outdoor temperature of 20°C, for an occupant arriving from a long absence, and indoor temperatures of 20°C and 25°C, the probabilities of opening a window are .007 and .033, respectively, meaning that statistically, an occupant would open his or her window in 0.7% or 3.3% of simulations at those indoor and outdoor conditions. If a window opening was predicted, then the duration of window opening is drawn from a Weibull distribution, and will remain open for the predicted period unless conditions cause an occupant to close the window prematurely. Similar computations are conducted for window openings and closings during intermediate presence and upon departure.

In the RSF building, two of every three windows on the southern façade are manually operable, and the third is automatically controlled. Windows on the north façade are divided between upper- and lower-windows; every other upper window is automatically controlled, while all of the lower windows are manually controlled. In the single-zone model, there are six windows on the southern façade (Figure 5.5, and twelve on the northern façade; six lower and six upper. In total there are ten manually controlled windows, five automatically controlled windows, and three fixed windows. The manually operable windows are four on the southern façade and six on the northern façade; the automatically controlled windows are two on the southern façade, and three upper windows on the northern façade. Partial window openings are not considered, in both the real building and in the building model, all openable windows are either fully open or fully closed at all times, and when fully open, windows act as connections between nodes in a nodal air flow network as described by Gu et al. in [34].

5.3.2.2 Occupant Presence

The second model of occupant behavior applied in this study is the occupancy prediction algorithm developed in [75]. This model takes the form of an inhomogeneous Markov chain, and uses daily probability profiles to determine occupancy at 15-minute intervals for a full year for each occupant in a given zone. The model accounts for both short and long absences (i.e. lunch-breaks and multi-day absences), and provides a critical input for other stochastic occupant behavior models. Details of the implementation of this model are included in Section 5.2.1.

Many behavior models account for occupant presence, and cite higher probabilities of occupant actions when occupants enter or leave a space [50, 83, 94]. The Page occupancy model provides more granular information on when occupants are arriving and departing, and whether their absences and presences are long or short; information which feeds into the window-use model. If a standard schedule of occupancy were used instead, there would be only three cases of occupancy: arriving in the morning after a long absence, present all day, or departing at the end of the day for a long absence. With the more granular presence profile, we can have occupants arriving, departing, or present for long or short time periods (six different cases), spread out appropriately through each day and each week.

One detail of the implementation is that the random seed for each simulation must be carefully controlled, since it governs the stream of random numbers which in turn governs stochastic occupant behavior within the model. The random seed can be set in one of two ways, either with a unique value for each simulation (to evaluate the impact of occupant behavior), or with the same value in each simulation for comparing multiple cases where some **other** parameter (e.g. fan control logic) is changed.

5.3.3 Generalizing Behavioral Models

In all of the work presented here, occupant behavioral models (OBMs) are used without modification, and assumed to be applicable to the SMM1 and the RSF building models used in optimization studies. In the case of the more generic SMM1 building model, it would be hard to argue that the OBMs are not applicable to the building model. In the case of the RSF building, or any real building for that matter, one can argue that the OBMs ought to be adapted in some way to the specific physical building being studied, or that the results of simulations might be skewed in some way. This section provides a discussion of this topic in general, with a focus on the model of manual window operation and the RSF building which are both the primary models in this research.

5.3.3.1 Consider Manual Window Models

First consider what the potential impacts of manual operation of windows on building performance (energy consumption and comfort) might be. Opening of a manual window can influence energy consumption in three primary ways, by changing the heating load, the cooling load, or the ventilation requirement for a space. If a window-HVAC interlock system is in place, any instance of opening a manual window will reduce the mechanical ventilation required. Depending on the season and the HVAC controls in the building, opening a window when outdoor conditions are cool will drive the cooling load down, or the heating load up, and opening windows when outdoor conditions are relatively warm will have the opposite effect on heating and cooling loads. Finally, opening windows when outdoor humidity is high can increase the latent load experienced by the building's HVAC system.

Second, recall that people tend to adjust their environment to make themselves more comfortable. When a cool breeze is coming in the window and we feel too cold, we might close the window, turn up the heat, or put on a sweater. In a perfectly conditioned and ventilated building, a perfect occupant would theoretically never need to adjust any part of his or her environment. By now we know that occupants are not perfect, but are somewhat predictable, and perfectly conditioned buildings are an impossibility. Just consider the Fanger PMV model for comfort, which states that there will always be at least 5% of people unsatisfied, regardless of indoor conditions.

Third, consider the Haldi model for manual window operation. The model captures both the predictable occupant behavior - when they open windows to get some natural cooling, for example, and the unpredictable occupant behavior - when they are unlikely to open or close a window according to predicted probabilities, but **do so anyway** thanks to the stochastic nature of the model.

5.3.3.2 Transferability

Now, examine in detail the differences between the LESO-PB building, where data was collected to develop Haldi's model, and the RSF building, where the model is used to simulate occupant behavior. Table 5.1 highlights the differences between the two buildings, which are significant. To account for the difference in office layout (single occupancy vs. open plan), we assume that one occupant has 'ownership' of each manually operable window in the open plan spaces in the RSF building, which from experience, is true. Generally there is one occupant closest to each window and takes responsibility for opening and closing it.

Table 5.1: Differences between the LESO and RSF buildings.

Characteristic	LESO	RSF
Office layout	Perimeter dual and single occupancy offices.	Primarily open plan, with some perimeter single offices.
Window open/close notification system	None.	Popup notification on computer desktop.
Heating Cooling Ventilation	Radiant Heating No Mechanical Cooling No Mechanical Ventilation	Radiant Heating Radiant Cooling UFAD
Window arrangement	Each occupant has an individual view window.	Some occupants have their own window; some share access to open plan windows.
Window Hinge	Bottom and side hinge.	Top hinge only.
Climate	Dry / cold in winter Mild / hot in summer HDD(65F) 5221 CDD(65F) 584	Dry / cold in winter Dry / hot in summer HDD(65F) 8010 CDD(65F) 278

Other differences between the two buildings are more difficult to account for. The presence of a notification system in the RSF might mean that occupants are **much** more likely to open or close a window when receiving a notification than at other times. This feature is not included in Haldi's model, nor is there a straightforward way to incorporate an additional tuning parameter without first conducting a field study to understand how occupants in the RSF do or do not respond to the notifications.

The differences in climate and in HVAC system are potentially the most troubling. The lack of any ventilation system in the LESO building means that LESO occupants are probably much more likely to open windows for fresh air than RSF occupants. The lack of any cooling system in the LESO building appears to have the same effect at first, and one might think that LESO occupants would open windows more for cooling than RSF occupants. Note that the OBM for manual windows is dominated by indoor temperature, and the probability of an occupant opening a window is high when indoor temperature is high (other parameters held constant). Given that the RSF building is cooled, RSF indoor temperatures are typically so low that the probability of opening a window is under 0.1, according to Haldi's model. So where a LESO occupant might open windows often for cooling in the LESO building, he or she might never feel the need to open a window in the RSF, simply because the RSF is mechanically cooled.

5.3.3.3 Adapting OBMs

Truly, there is no better way to adapt a behavioral model from one building, region, or community to another is to conduct a field survey, and to tease out differences in model parameters that vary according to some rules, i.e. Europeans open windows according to X, and Americans open windows according to Y, or model coefficient A can be scaled linearly according to latitude. It may also be that there are not significant differences, and that a single model is highly transferable; without many more years and buildings of data to work with, it is impossible to say.

Considering the RSF and LESO buildings, a few thoughts on adapting Haldi's model to the RSF population are to:

(1) Add a scaling factor that increases the probability of opening or closing a window whenever occupants receive a notification to do so.

(2) Adjust indoor temperature model coefficients to account for the tighter temperature control in the RSF building.

In **(2)** above, the intent is to account for the fact that the LESO building is not cooled, so occupants experience a much wider range of indoor conditions, and it follows that they are more willing to accept those conditions. Since RSF occupants experience a very narrow band of indoor conditions, it may be that they are more likely to open a window when indoor conditions deviate even a little bit from typical. The coefficients in the model that correspond to indoor temperature could be scaled up to account for this difference.

5.3.3.4 Impact on Results

With the discussion of differences between the RSF and LESO building in mind, and the fact that the model was used without modification or adaptation, what is the likely impact on results of simulation studies?

If anything, the results of simulations with the RSF model are probably conservative. This is due to the fact that the main driver (after arrival/departure status) for manual operation of windows in the behavioral model is indoor temperature. The LESO building uses passive night ventilation (occupants can leave their windows open at night) for cooling, and lacks any mechanical cooling, so temperatures can vary significantly. The RSF building however is very tightly controlled by a massive and responsive embedded radiant system. Thus, conditions in the RSF are very rarely conducive to opening and closing windows, according to Haldi's model of window operation - and the impact of occupant window use will be small.

What about the window open/close notification system in the RSF? This probably leads to occupants doing the right[1] thing with their windows more often than not. The notification system is tied directly to the RSF automatic windows, and instructs occupants to open or close their windows whenever the automatic windows do the same. When the RSF building was designed, the natural ventilation system was designed to be able to achieve sufficient airflow (for night cooling) without the help of occupants (manual windows). That is, the openable surface area of the automatic windows is sufficient, and the openable area of manual windows is essentially a bonus feature for occupants.

Recall that the RSF automatic window controls only allow windows to be open when outdoor conditions are virtually identical to indoor conditions, so there is no significant change in heating or cooling load, and the RSF mechanical ventilation system operates independent of the natural ventilation system, so there is no impact on mechanical ventilation. If in fact occupants open and close their windows whenever the notification system instructs them to, the impact on building energy consumption is negligible. Rather, it is those occasions when occupants open or close windows **outside** of the more predictable times that have a greater impact on building performance. The instances when occupant use of windows can impact performance are those instances when they leave a window open on a hot afternoon, or overnight when it is very cold out.

The results of optimization studies presented later often show the average of many simulations, or the distribution of many simulations. If in fact Haldi's model of occupant window operation is not applicable to the RSF, then it might mean the average results are inaccurate, or that the distribution of results is inaccurate.

It is entirely speculative, but the opinion of this author is that the average results would not change significantly even if there were large differences in a behavioral models. Performance of the RSF building is not very sensitive to occupant behavior. Certainly there are isolated events where a number of occupants all do the same thing to improve or degrade

[1] Assuming that the notification system is telling occupants the right thing.

performance, but over a longer timescale than a few days, the building's systems, controls, and the outdoor conditions have a much higher impact on performance than occupant behavior. The distribution of results however, is more challenging even to speculate on. In general, distributions of energy consumption show a long right tail, meaning that occupant behavior can do more harm than good. Occupant behavior can lead to some energy savings, **or** some increased energy consumption, but relatively speaking, it can lead to larger energy increases than decreases. This makes sense, since there is some lower limit of energy consumption for a building, but theoretically no upper limit, so the benefits of occupant behavior hit a lower bound but not an upper bound.

Chapter 6

Occupant Behavior in Mixed Mode Buildings

6.1 Introduction

In this chapter, building energy simulation models are coupled with stochastic occupant behavior models, and the magnitude and distribution of impacts that occupants have on building energy consumption are demonstrated via a simulation study. The results show that occupant actions can increase or decrease energy consumption depending on the HVAC control strategy implemented. For a single month during the cooling season, the range of HVAC electricity consumption predicted by a set of simulations that included stochastic models of occupant window, blind, and lighting use varied by approximately 20% for each of 25 different control scenarios.

6.2 Results

Figure 6.1 shows the results of 400 simulations using the building model described above coupled with the model for occupant use of manual windows; it shows the variation of building HVAC energy consumption for the month of June including the effects of manual window use. In this case, the deterministic simulation (without occupant behavior models - all occupant windows closed) resulted in an energy consumption of 4581 kWh, and the average result of the stochastic simulation (with behavioral models) is 4882 kWh. Note that the resulting distribution of energy consumption is skewed to the right, indicating that the response of the building to occupant actions is not normally distributed, and that occupant action can lead to cost penalties or cost savings compared to the deterministic case (or to the average result of the stochastic case). The long right tail on the data shows that occupants can incur larger energy consumption penalties than they can savings; if the distribution had a long left tail, it would indicate that the opportunity for savings from occupant interaction is much greater than the risk of penalties. While the 6% difference between deterministic and stochastic results appears small, the second value comes with richer information: a lower and upper limit defining a range of what energy consumption could be, and most importantly a high level of confidence in the average result. A generalized extreme value distribution was fitted to the data; the shape k, scale σ, and location μ parameters for the distribution are given in Table 6.1.

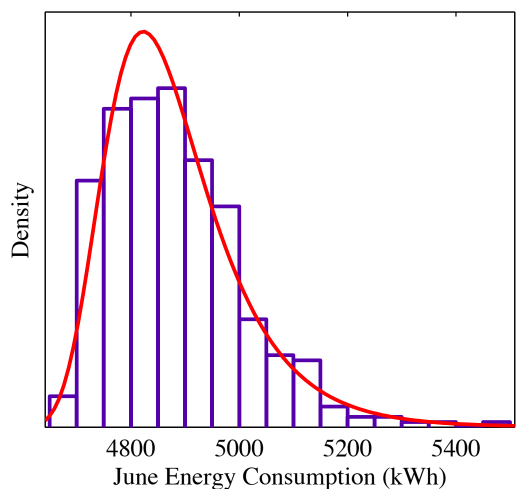

Figure 6.1: Distribution of energy consumption for 400 samples of window opening behavior.

Figure 6.2 shows data similar to that in Figure 6.1 for 25 unique BAS-window control scenarios to demonstrate the dynamic response of the building and occupants to different building control methods. Recall that the building energy model has two banks of windows; the lower bank is assumed to be occupant-controlled (OCC), and the upper bank is controlled by the building automation system (BAS). For simplicity, the various BAS window control scenarios are defined as a single window opening event with two defining parameters: opening

Table 6.1: GEV fit parameters.

Parameter	Estimate	Std. Err.
k	0.0214	0.0403
σ	95.52	4.039
μ	4825	5.449

time and duration; this entails a single control action that is repeated by the BAS on each day for the duration of the simulation. Five unique opening times (8, 9, 10, 11, and 12:00) and five opening durations (1-5 hrs) represent the assumed set of possible BAS-window control options, resulting in a total of 25 options. For each control option, 400 simulations were completed, so the results presented in Figure 6.2 are the product of 400x25 = 10,000 simulations.

Note the location of the vertical line in Figure 6.2 (indicating the result of a deterministic simulation) relative to the distribution of results from stochastic simulations; the bottom row of plots for examples shows that when BAS windows are opened for one hour, occupant actions lead to higher energy consumption than the deterministic simulation predicts.

When the BAS window signal commands windows to open at 11:00 or 12:00 for 4 or 5 hours (top-right), occupant behavior leads to lower simulated energy consumption than what is predicted by a deterministic simulation without occupant models. In contrast, in the lower-left quadrant of the figure, the deterministic result is significantly lower than the results predicted by the simulations that incorporate stochastic behavior. For this window control scenario depicted in Figure 6.2, while the stochastic cases reveal significant distributions of energy consumption for each individual control option, the range of energy consumption between the minimum and maximum mean (5.0-5.7 MWh) is smaller than the range of energy consumption of the deterministic case (4.6-6.0 MWh). For this analysis, occupant behavior compresses the range of building energy use, dampening the effects of the BAS window controls. Looking at the figure as a whole, note the general trend from left to right and from bottom to top (opening BAS windows for longer periods, later in the day) of an increase in

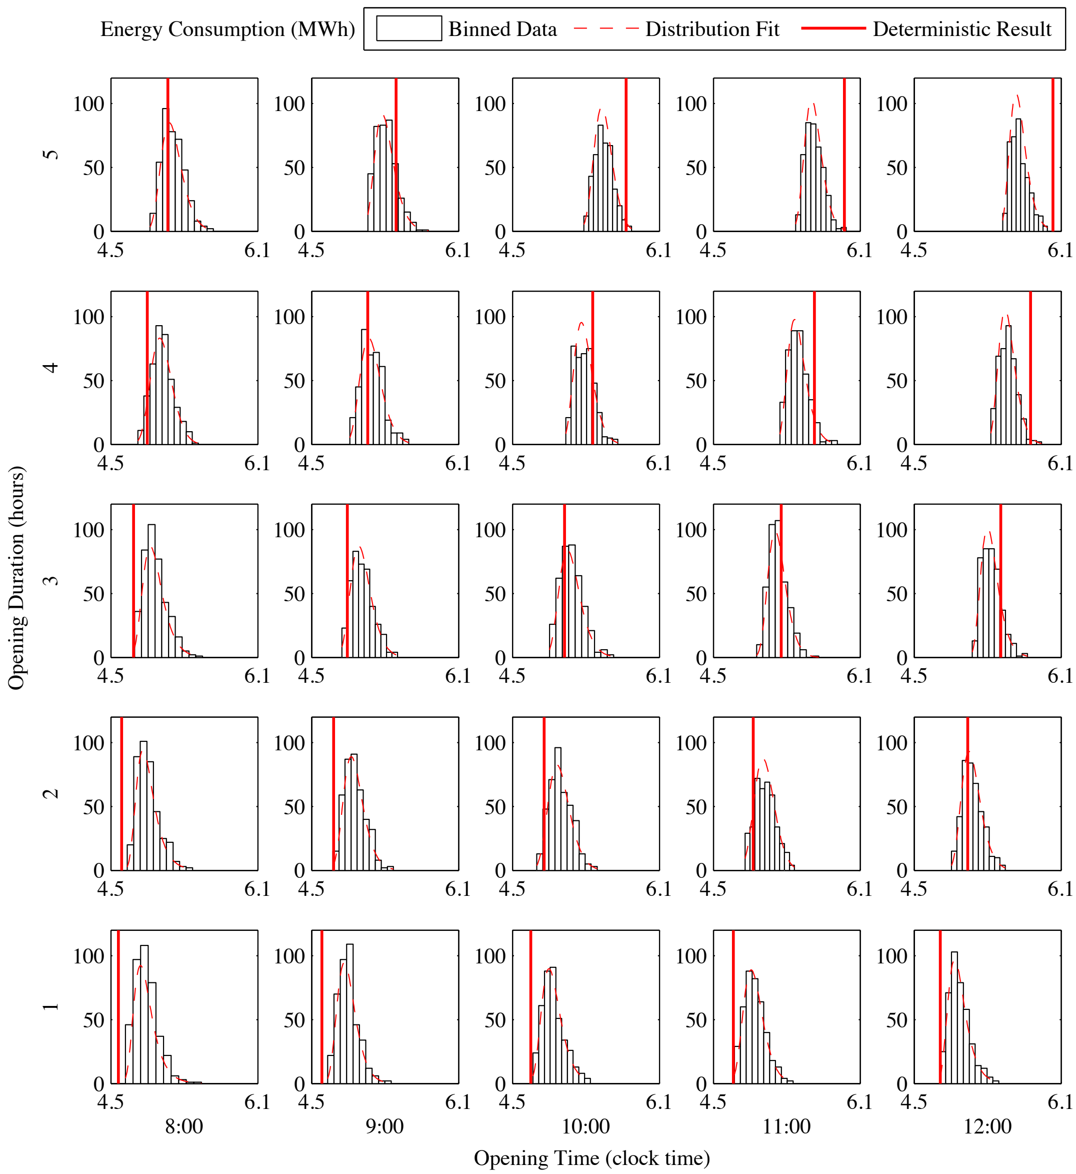

Figure 6.2: June energy consumption for 25 control options.

energy consumption. This makes sense in the month of June, when opening windows in the afternoon may allow warmer air into the building, thereby increasing cooling loads. In the central plot in Figure 6.2, there is good agreement between the average stochastic result and the single deterministic result, which is strictly coincidental; for the particular time period of the simulation and the building and occupant behavior models used in this study, when BAS windows are opened at 10:00 for 3 hours, occupant behavior does not shift average energy consumption up or down from the deterministic baseline.

Figure 6.3 shows the effects of including additional behavioral models in a suite of simulations just as in Figure 6.2, however here four different building cases were included: a deterministic case (D), a case which includes the stochastic window use model (W), a case that includes use of windows and blinds (WB), and a case that includes models for window, blind, and lighting use (WBL). Each of the four cases is identical to the others with the exception of the set of behavioral models included.

Two observations are made: First, in all cases, the inclusion of more behavioral models leads to lower energy consumption; the WBL cases use less energy than the WB cases, which use less energy than the W cases. Second, in some cases the deterministic energy consumption is significantly lower, in others significantly higher, and yet in others it is similar to the occupant influenced scenarios, so it can not be assumed that the inclusion of stochastic behavior always leads to decreased energy consumption when compared to deterministic simulations.

Where Figure 6.3 looks at three combinations of occupant behavior models incorporated in the building model, Figure 6.4 includes all 8 permutations of the three behavioral models. Clearly those that include window opening behavior have the largest impact, which is to be expected since we consider only the impacts on HVAC performance. Again we see that the impact of occupant behavior is different with each different automatic control; the lower-left control cases in Figure 6.4 seem to result in consistent performance regardless occupant behavior, and occupant behavior in these cases has more potential to increase energy

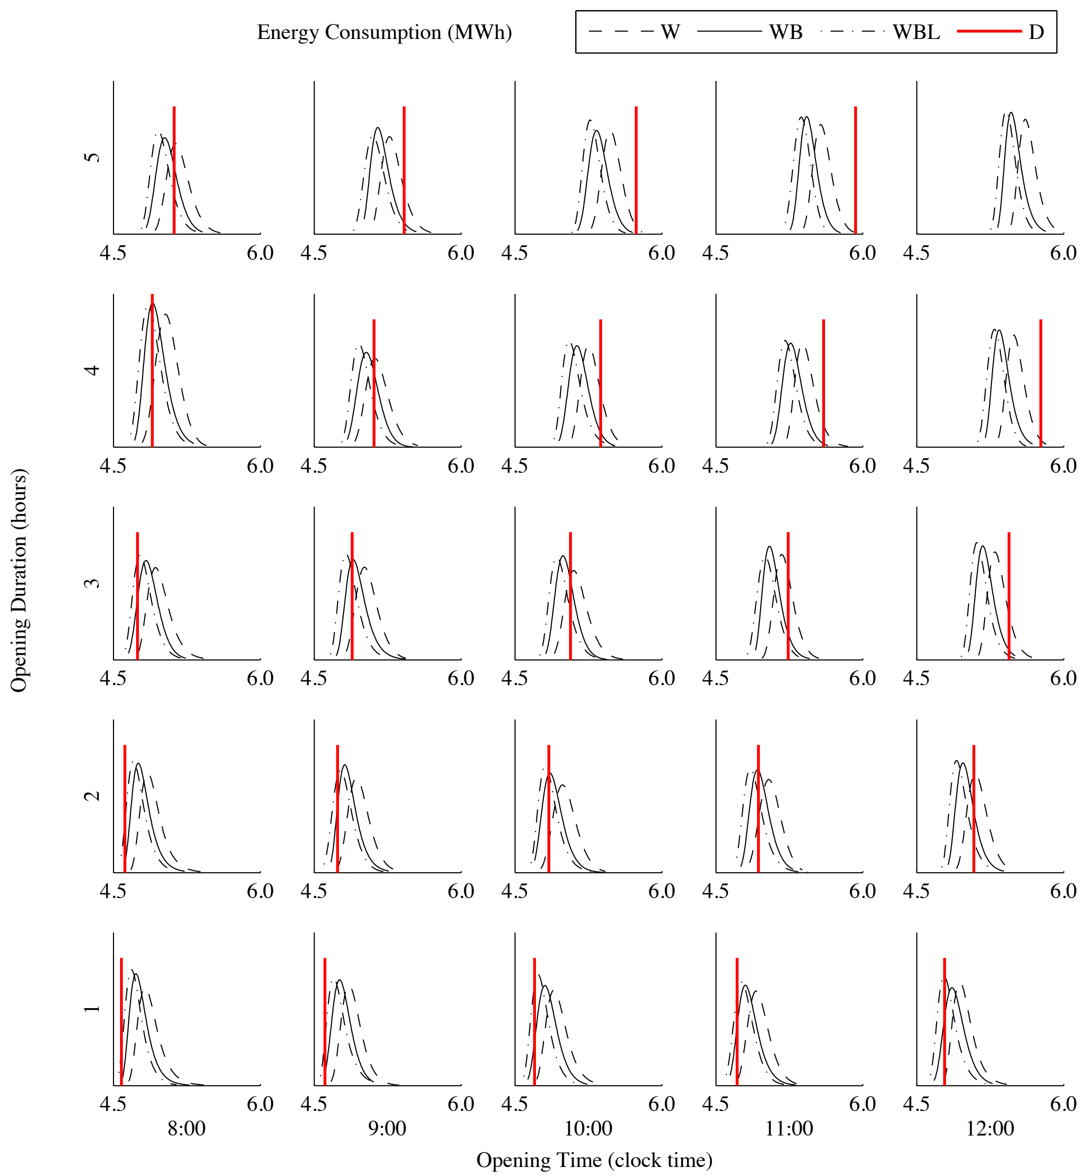

Figure 6.3: PDF of building energy consumption with three models of occupant behavior.

consumption than to decrease energy consumption. The automatic control cases in the top right however are much less robust to the impact of occupant behavior; when models of occupant window use are included, HVAC energy consumption is roughly 10% lower than when models of occupant window use are not included.

Given the results presented in Figure 6.1, Figure 6.2, and Figure 6.4, it is clearly difficult to predict exactly what the net effect of occupant behavior in a building will be, especially for a range of different control scenarios issued from a BAS. This fact leads us to propose the development of methodology for determining near-optimal control rules for buildings that are sensitive to occupant interaction. Work currently underway is thus concerned with developing such an optimization environment.

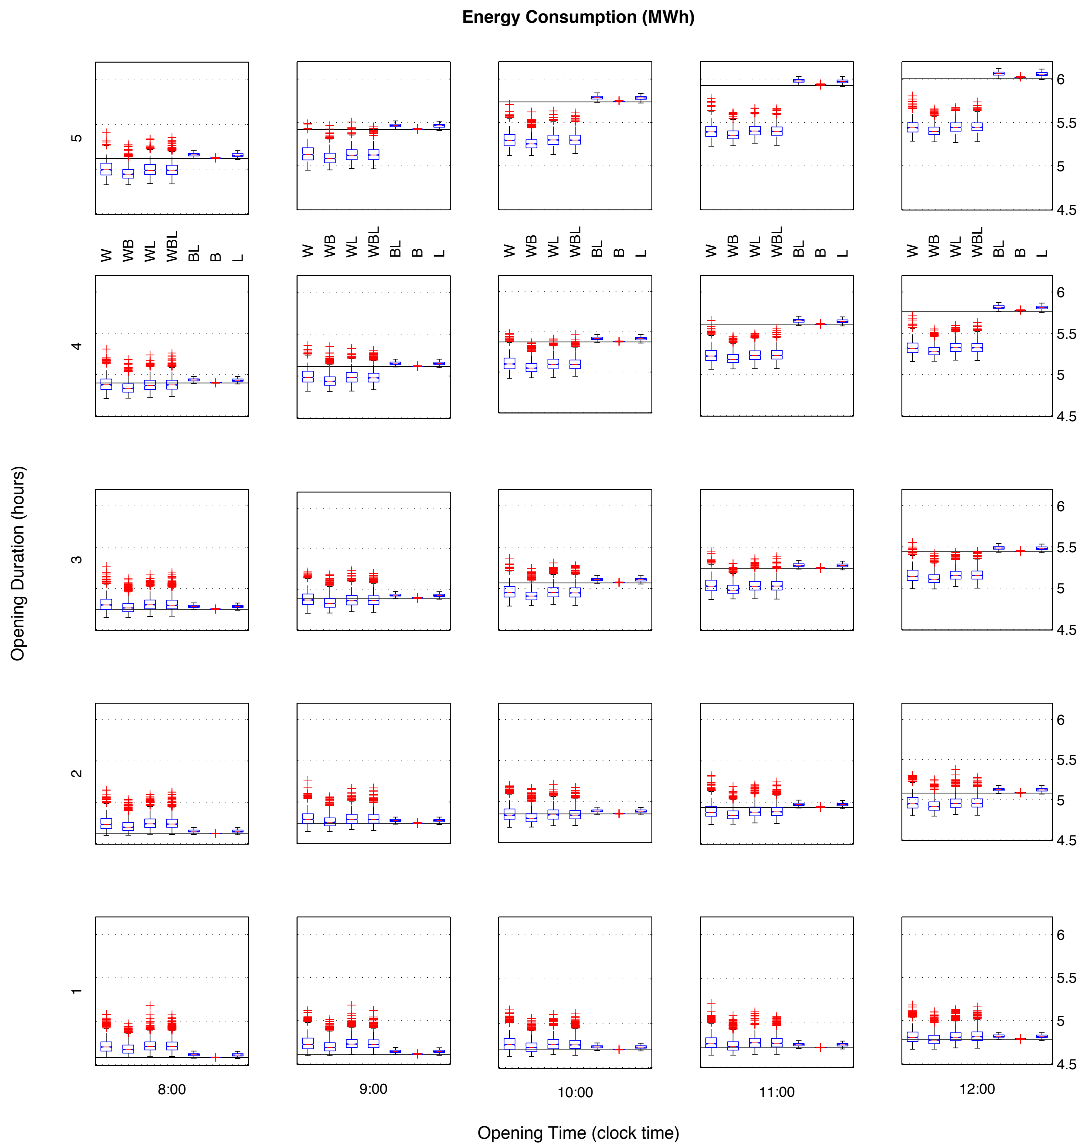

Figure 6.4: Energy consumption for 25 control options considering occupant use of windows (W) blinds (B), and lights (L), in various combinations. The solid horizontal line corresponds to the result of simulating with no occupant behavior.

6.3 Conclusion

An investigation of the effect of occupant behavior in buildings on building HVAC electricity consumption was conducted. The results show that occupant actions can increase or decrease energy consumption depending on the HVAC control strategy implemented. Additionally, it is shown that as more models of occupant behavior are added to the energy model, the predicted energy consumption tends to decrease, indicating that added access to manually operable systems and occupant actions on those systems are likely to reduce energy consumption. For a single month during the cooling season, the range of HVAC electricity consumption predicted by a set of simulations that included stochastic models of occupant window, blind, and lighting use varied by approximately 20% for each of 25 different control scenarios.

Chapter 7

Bottom Up Results: Tuning Setpoints

7.1 Introduction

In the context of optimizing controls for high performance buildings to achieve energy savings while maintaining or improving occupant comfort, this study is aimed at finding incremental improvements in the existing control logic for an existing building, specifically the RSF Building. Through the use of parametric studies with a building energy simulation model that incorporates a stochastic model of occupant window use and automatic window control logic commensurate with the physical building, the performance of different natural ventilation control strategies is evaluated. Due to the inclusion of stochastic occupant behavior in the building model, results of simulations are presented as a distribution of potential outcomes, rather than a scalar value, however the average of these distributions is used for simpler comparison between control strategies. The results show that annual HVAC electricity consumption can be reduced by 3% to 15%, depending upon the controls changes implemented.

In the following the building selected for the study and a simplified energy model for that building are described. Special attention is given to the HVAC and automatic window controls in the building, and the parametric studies conducted are presented.

7.1.1 Building HVAC and Window Controls

The main goal of the parametric study is to determine if savings can be achieved through decreased demand on HVAC equipment, specifically the air distribution and radiant systems. The radiant system is controlled to maintain a weighted average space temperature, where dry-bulb temperature contributes to 80% of the average, and the radiant slab temperature contributes to 20% of the average. Supply water temperatures are set according to a seasonal reset schedule, and circulation pumps are cycled on whenever the space temperature falls outside of the range between the heating and cooling setpoints. The UFAD system is controlled to maintain adequate CO_2 levels in each space; a variable air volume ventilation system modulates between a minimum setting and a fully open setting as space-CO_2 levels change per ASHRAE Standard 62.1-2010 [5], and the outdoor-air fraction is set to the maximum required per person or per floor-area (whichever is greater).

7.2 Automatic Window Control Strategies

The RSF automatic window control strategy began as a simple night purging sequence, then evolved into a 24-hour natural ventilation scheme. A proposed change is the addition of an interlock system, which would turn off mechanical ventilation entirely to a given zone when natural ventilation is in effect in that zone; this represents a significant change from concurrent to changeover control as described above. The three window control scenarios investigated are shown graphically in 7.1. In the first control scheme, windows are opened at night when temperature and humidity are both low in order to cool the building and flush out stale air. In the second scheme, windows are opened at any time when outdoor conditions allow, and in the third (proposed) scheme, the ventilation system is turned down significantly when windows are opened.

In all three cases, the same temperature, wind, and humidity-based setpoints are used, and there is a lockout function that closes the windows whenever wind is gusting. A 1°F deadband exists around all temperature setpoints to prevent excessive open/close cycling of the automatic windows; the actual temperature setpoints are discussed further below. The humidity setpoints allow windows to open when outdoor relative humidity (RH) drops below 50%, and force windows to close when outdoor RH rises above 52%. The wind speed lockout control algorithm uses a combination of time and windspeed to determine when windows can be opened or closed. If wind speeds remain above 17 mph (7.6 m/s) for six minutes, the windows are forced to close, and cannot re-open until wind speeds have remained below 12 mph (5.4 m/s) for five minutes.

7.2.1 Control Heuristics

Considering the systems and controls in place, several straightforward methods exist for achieving the stated goal, which is to offset cooling, or ventilation requirements by employing natural

Free Cooling. To offset a cooling load, the windows need to be opened when the building requires cooling and the outdoor temperature is cooler than the indoor temperature.

Pre-Cooling. In a building like the RSF which features significant thermal mass, if a warm period is preceded by a cooler period, it is also possible to pre-cool the buildings thermal mass in anticipation of the warm period, potentially avoiding the need for cooling entirely during the warm period.

Natural Ventilation. In order to offset a ventilation requirement, automatic windows should be opened when outdoor air conditions fall within established comfort limits, i.e. per ASHRAE Standard 55-2010 [4], which provides a comfort temperature range that varies with the monthly mean outdoor air temperature.

If the air and slab temperatures are kept between the heating and cooling setpoints, the radiant system will not cycle on and thus will not consume energy. In this case, the radiant

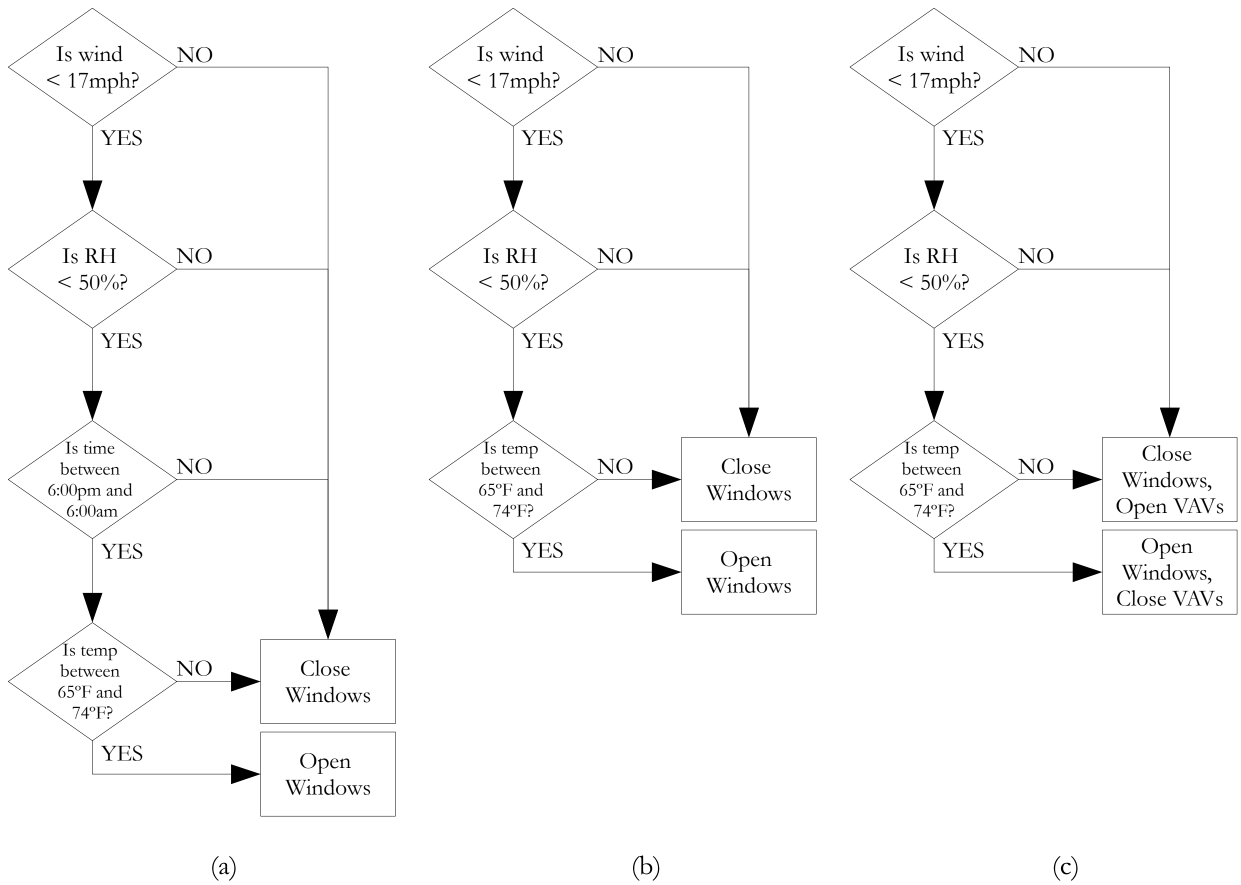

Figure 7.1: Diagrams for three control logic scenarios considered for parametric study and for implementation in the RSF building; the RH conditional statement refers to the relative humidity (RH) of outdoor air.

system responds to conditioning provided by natural ventilation without extra control logic, and savings will accrue without extra effort. In the case of ventilation however, the situation is different; the ventilation system will typically continue to provide fresh air (dampers set to a minimum or unoccupied setting), even when windows are opened. The controls of the mechanical ventilation system must be able to respond to the state of the windows and ramp down when windows are open, otherwise the system may continue to supply air when it is not needed, and potential energy savings will not be realized.

7.2.2 Parametric Studies

Three parametric studies are presented. In each case a small set of control parameters are chosen from the logic diagrams presented above, and reasonable values of each are chosen for further analysis. For each combination of values, a new building energy model is created with the unique set of values employed in the simulated control system. In the second study presented, two parameters and eight values of each are chosen, leading to 64 (8x8) combinations of control parameters, and 64 unique building energy models. Each unique model is then simulated 400 times, and in each simulation the simulated occupant behavior is slightly different due to its inherent stochastic nature. An investigation into the necessary number of samples to achieve convergence in the average energy consumption revealed that samples greater than 200 are sufficient, with negligible changes observed for larger samples. We adopted 400 samples where possible and reverted to 200 samples for time intensive investigations. The individual results of these 400 simulations are then combined and lead to a distribution of results, the average of which is taken to represent the entire set. The 64 average results are then compared, the result which minimizes energy consumption while preserving occupant comfort is chosen as the best, and the corresponding control parameters are recommended. **Night Purging.** In the night purging control logic (see Figure 7.1(a)), three setpoints were selected for a parametric study: the two time-of-day limits (6:00 pm and 6:00 am), and the lower temperature limit (65°F (18.3°C)). Due to the infeasibility of some combinations (i.e. start time equal to or later than end time), only 75 combinations of the selected control parameters are possible.

Table 7.1: Parametric study 1: night purge control parameter values

Night Ventilation Start Time	Night Ventilation End Time	Outdoor Air Temp. Lower Limit
18:00	21:00	61°F, 16.1°C
21:00	24:00:00	63°F, 17.2°C
24:00:00	3:00	65°F, 18.3°C
3:00	6:00	67°F, 19.1°C
6:00	9:00	69°F, 20.6°C

Natural Ventilation. In July 2012 the night-purge control logic was abandoned in favor of a 24-hour natural ventilation scheme, a new set of control parameters were selected for analysis. Here, the upper and lower temperature thresholds for window opening were adjusted.

Window-HVAC Interlock. The final control change investigated was a direct linkage between the automatic windows and the ventilation system (see Figure 7.1 (c)). Without the linkage, the ventilation system operates normally whether windows are opened or closed (concurrent control), and with the linkage, the local ventilation system is turned off when windows are open, and turned back on whenever windows are closed (changeover control). In EnergyPlus, this is achieved by setting the availability of terminal units to zero, while in reality it might be achieved by closing a VAV damper completely. With the existing CO2-based control logic, VAV dampers are set to minimum setting but not closed completely when CO2 levels are low; this ensures that there is always sufficient fresh air, whether windows are open or closed.

Table 7.2: Parametric study 2: natural ventilation control parameter values

Outdoor Air Temp. Lower Limit	Outdoor Air Temp. Upper Limit
62°F, 16.7°C	71°F, 21.7°C
63°F, 17.2°C	72°F, 22.2°C
64°F, 17.8°C	73°F, 22.8°C
65°F, 18.3°C	74°F, 23.3°C
66°F, 18.9°C	75°F, 23.9°C
67°F, 19.4°C	76°F, 24.4°C
68°F, 20.0°C	77°F, 25.0°C
69°F, 20.6°C	78°F, 25.6°C

7.3 Results

In the following discussion, the results of the night purge control logic cases are presented in Figure 7.2 through Figure 7.4; Figure 7.2 provides a comparison of the results of all 75 combinations of control parameters considered, while Figure 7.3 and Figure 7.4 show the energy performance and comfort performance for the default case and for the 8 alternatives that resulted in the lowest energy consumption. Results of the second parametric study, which considers 24-hour natural ventilation, are shown in Figure 7.6; in this figure, energy consumption resulting from each combination of control parameters is displayed as a function of each month. Figure 7.7 shows monthly energy savings for the third parametric study, which explores the application of a window-HVAC interlock system. In all of the results presented here, energy consumption and comfort values are for the simplified 1-zone model, which has a 3600 ft2 floor area, and 18 occupants when fully occupied. Additionally, the results are often compared to what we have named the default case, which is the case with existing in-use setpoints; e.g. in the night purge control logic, the default case refers to the case with setpoints matching actual RSF night purge control setpoints at the time of the study; for each parametric study, there is a different default case, and thus a different benchmark for comparison.

In Figure 7.2, 75 values are presented; each value is the average HVAC electrical energy consumption (sum of fan, pump, and chiller electrical consumption) for 400 simulations of the single-zone RSF model for the month of June. Each value corresponds to a unique set of control parameters; for example, the circle corresponding to the existing control parameters represents 973 kWh, which is the average of 400 simulations with the existing control parameters: 6:00 pm start time, 6:00 am end time, and 65°F (18.3°C) minimum outdoor air temperature. Any circle that appears darker than this one corresponds to a combination of control parameters that lead to lower building energy consumption. Eight of the 75 simulations stand out as good performers, but no single combination leads to a dominant best. This graphic only shows the energy consumption of each case, and does not consider comfort conditions, however the modeled building does maintain similar comfort conditions in each case, as shown below in Figure 7.4. Several trends in the data presented in Figure 7.2 are evident; in the upper-right portion of the figure (early morning-ventilation), the 65°F setpoint is better than any higher or lower value, and the energy consumption for these cases are generally lower than those with night ventilation start times before midnight. This suggests that opening the windows prior to midnight does little to reduce energy consumption, however the case that resulted in the lowest overall energy consumption (6:00 pm start time, 6:00 am end time, 61°F (16.1°C) temperature setpoint) does include these evening window-opening actions. In the next two figures, the results of the 8 cases that yielded the lowest average energy consumption are compared to the default case, and a summary of each case is given in Table 7.3. Figure 7.3 shows the distribution of June HVAC energy consumption from 400 simulations, and Figure 7.4 shows a comfort metric, which is described below. In both cases, lower values of the data displayed (energy consumption or comfort) are better.

The comfort metric displayed in Figure 7.4 is intended to quantify the total comfort (or discomfort) for a given building and time period, similar to a total energy consumption metric. Here we sum the product of occupancy and the absolute value of the Predicted Mean Vote (PMV) comfort score as reported by the Fanger Comfort Model for each hour of

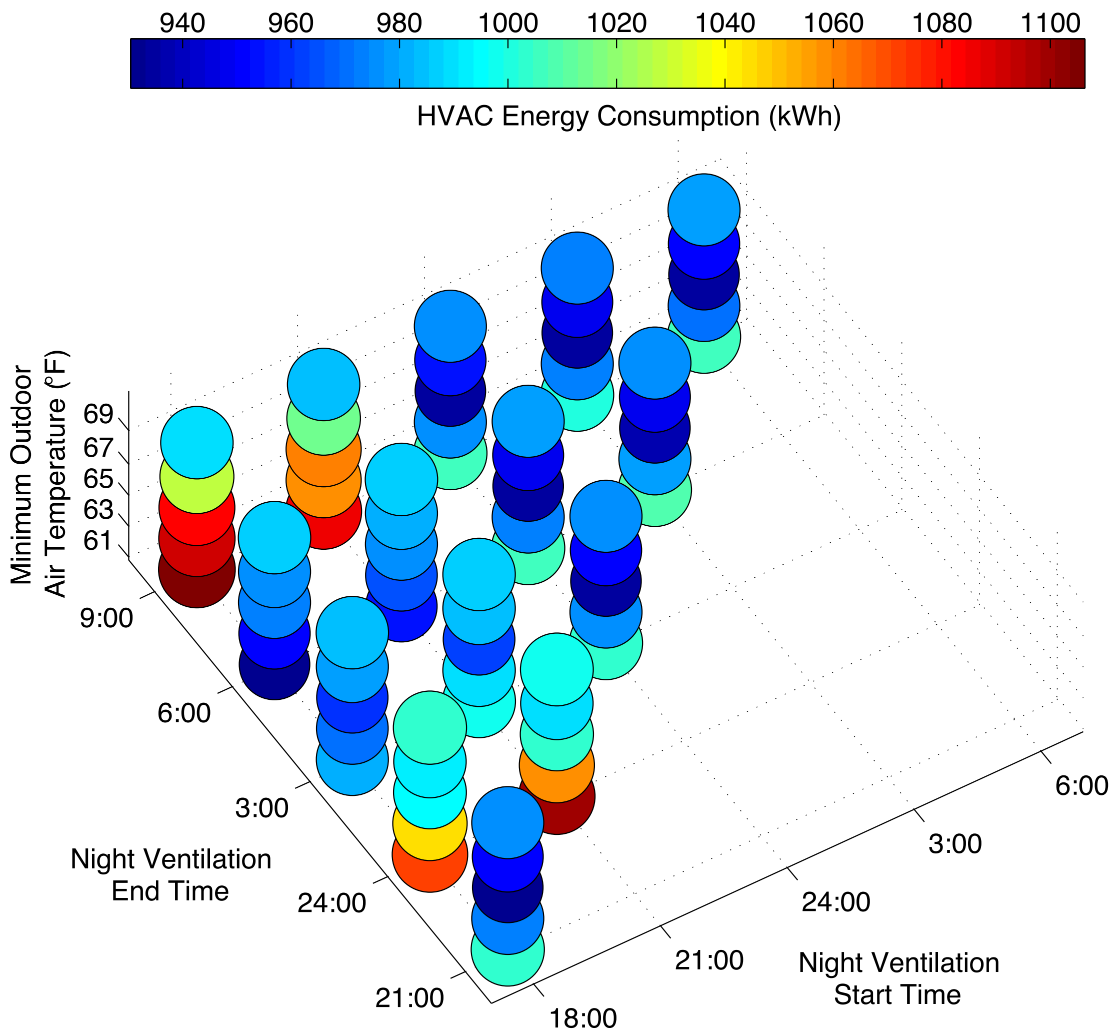

Figure 7.2: Energy consumption for the month of June for 75 combinations of night purge control setpoints. The default setpoints are 6:00 pm (start), 6:00 am (end), and 65°F (18.3°C) (outdoor air temp).

simulation according to Equation 7.1.

$$C = \sum_{t=1}^{T} |PMV_t| \, N_{Occ_t} \tag{7.1}$$

In Equation 7.1, C is the comfort metric value, t is an index for each time step, T is the final time index, PMV_t is the PMV value for each time step, and N_{Occ} is the number of occupants at each time step. The resulting number is a time- and occupant-weighted value that is only useful for comparing nearly identical simulations; the value has no meaningful interpretation by itself. Smaller values of C are better as they indicate closer proximity to comfort neutrality.

In Figure 7.3, more detail on the results of each case are presented, highlighting the influence of occupant behavior on energy consumption. Instead of a single average value for the suite of 400 simulations, we see the distribution of results for each case, with the average value for each case indicated by a dashed line. In both the default case and the result in Figure 7.3(b), there is a very strong modal result around the average, indicating that the corresponding combinations of control parameters are less strongly influenced by occupant behavior. In the remaining seven subfigures, the average value is slightly to the right of values that occurred with higher frequency, indicating that the average is skewed by high outliers, and may slightly over-estimate the most likely response.

Table 7.3 shows a summary of the values presented in Figure 7.2 and Figure 7.3; the default set of data is boldfaced. Note that the range in HVAC electric consumption is 4.4%, while the range in the comfort metric is 0.5%, indicating that the opportunity for savings comes with minimal risk of adversely affecting comfort.

In Figure 7.5, 64 values are presented for suites of annual simulations with and without HVAC-window interlock control logic; just as in Figure 7.2, each of these 64 values is the average result of many unique simulations, 200 for this set; each value also corresponds to a combination of two control parameters: the minimum and maximum outdoor air temperature setpoints between which automatic windows are allowed to open. The default control

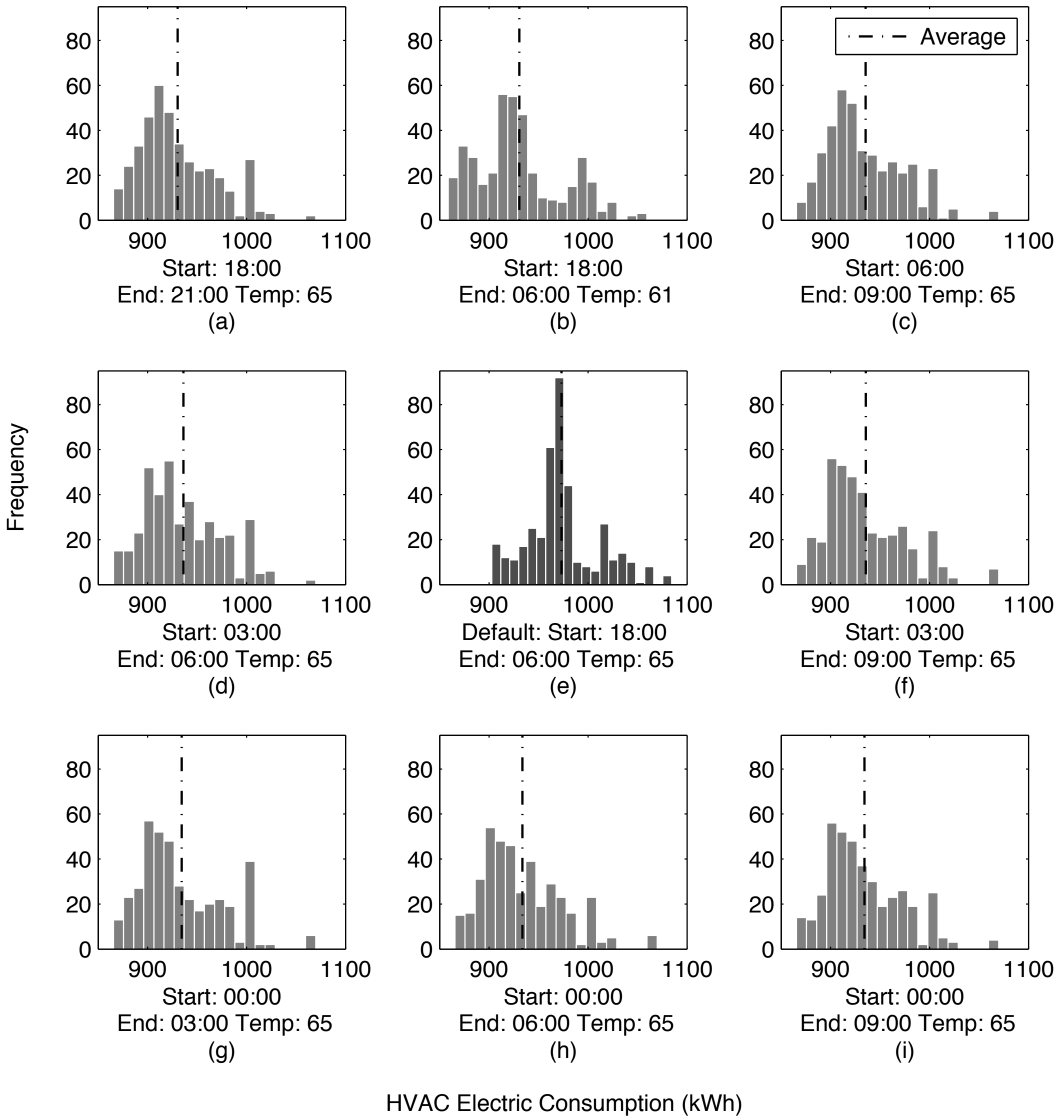

Figure 7.3: Energy consumption for the month of June for 8 combinations of night purge control setpoints, and the default setpoints (center).

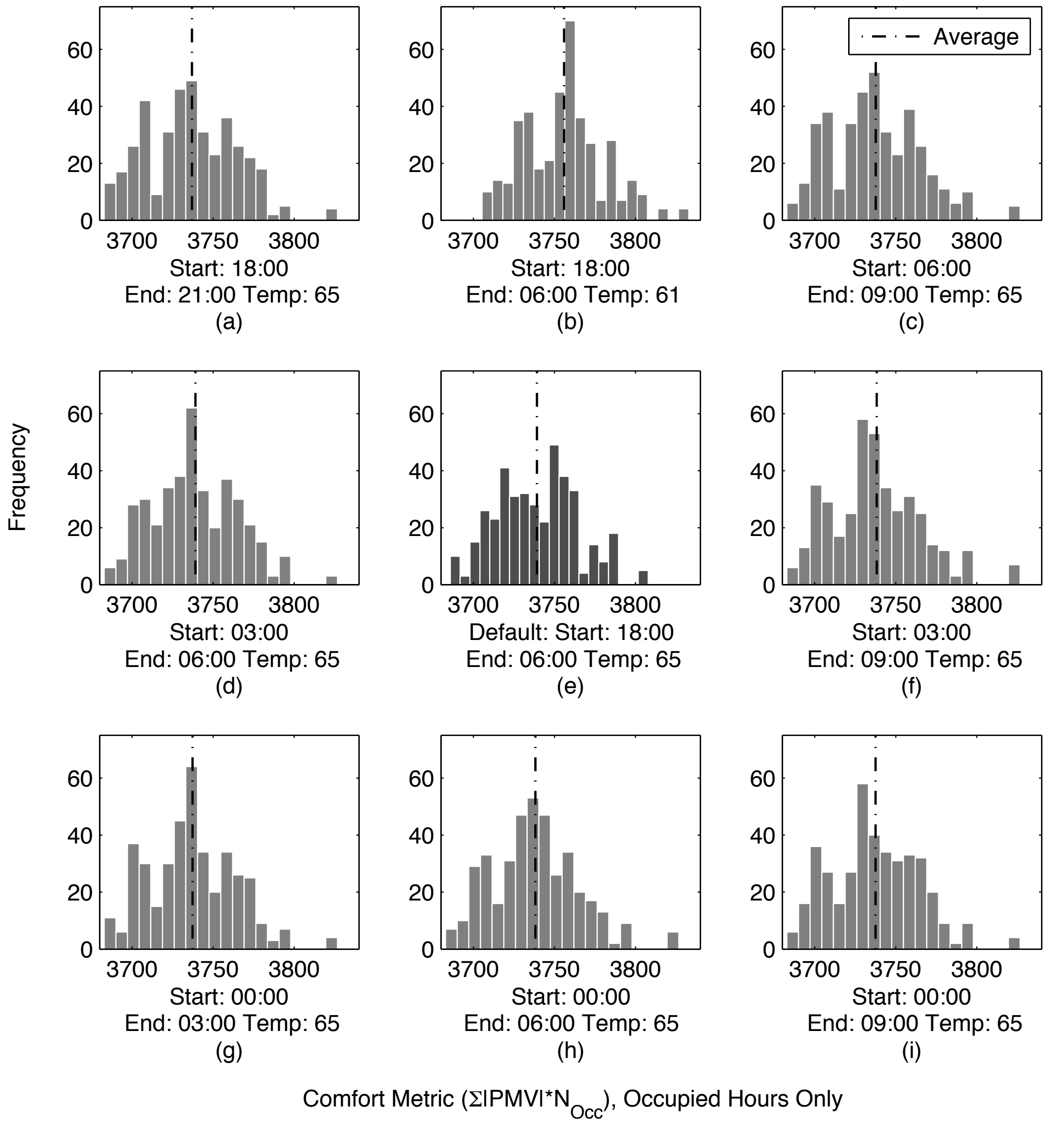

Figure 7.4: Cumulative comfort metric for the default case and for 8 cases that consumed the least energy on average.

Table 7.3: Summary statistics for the night purge control study.

$T_{OA,Max}$ °F, °C	Default-Case Energy Consumption (kWh)	Best-Case Energy Consumption (kWh)	Absolute Savings (kWh)	Percentage Savings
74, 23.3	191	189	2	1%
71, 21.7	147	142	5	4%
78, 25.6	306	285	22	7%
78, 25.6	333	253	80	24%
76, 24.4	404	355	49	12%
77, 25.0	747	723	24	3%
72, 22.2	961	942	19	2%
74, 23.3	1205	1142	63	5%
75, 23.9	381	343	38	10%
73, 22.8	213	199	14	7%
76, 24.4	318	245	73	23%
78, 25.6	142	137	5	4%

parameters are 65 and 74°F (23.3°C), and a white border highlights the corresponding value of energy consumption. Note that the energy values have been normalized such that the default value is equal to 1.0, so that the graphics show relative increases or decreases in energy consumption for different combinations of control parameters. For an automatic window control system without ventilation-interlock, savings are only realized when loads on the radiant system are reduced, as the ventilation system continues to operate regardless of window position. For a system with ventilation-interlock, savings accrue whenever windows are opened because fan energy and associated supply-air cooling loads are reduced. This difference between the two systems leads to considerably different choices of minimum and maximum outdoor air temperature for opening automatic windows. For a system without interlock, the combination of min- and max- temperatures that minimizes HVAC electric consumption is 67 and 74°F (19.4 and 23.3°C), while for a system with interlock the combination is 62 and 77°F (16.7 and 25.0°C). At this point it is meaningful to note that heating energy is not considered in this analysis, as electrical load reduction is the primary goal. Decreasing the setpoint for window opening will naturally lead to cooler indoor air temperatures when windows are opened on cooler days, and decreasing it to a value below the heating setpoint will cause the heating system to respond, and increase the heating energy required. The question then becomes, how much more costly is extra heating energy relative to electrical energy. To determine potential heating energy savings, a source energy analysis involving the wood pellet boiler and electricity mix of site generated and grid sourced electricity would have to be conducted, which is reserved for future work. Comfort is also treated explicitly in this analysis, as it was shown in the first parametric study that comfort conditions are consistently maintained by the HVAC system regardless of the window control logic. Note also that adverse comfort conditions are included indirectly through increased HVAC loads and energy consumption.

In Figure 7.6, 64 values are presented for each month of the case without an interlock system; just as in Figure 7.5, the values in Figure 7.6 have been normalized for easier

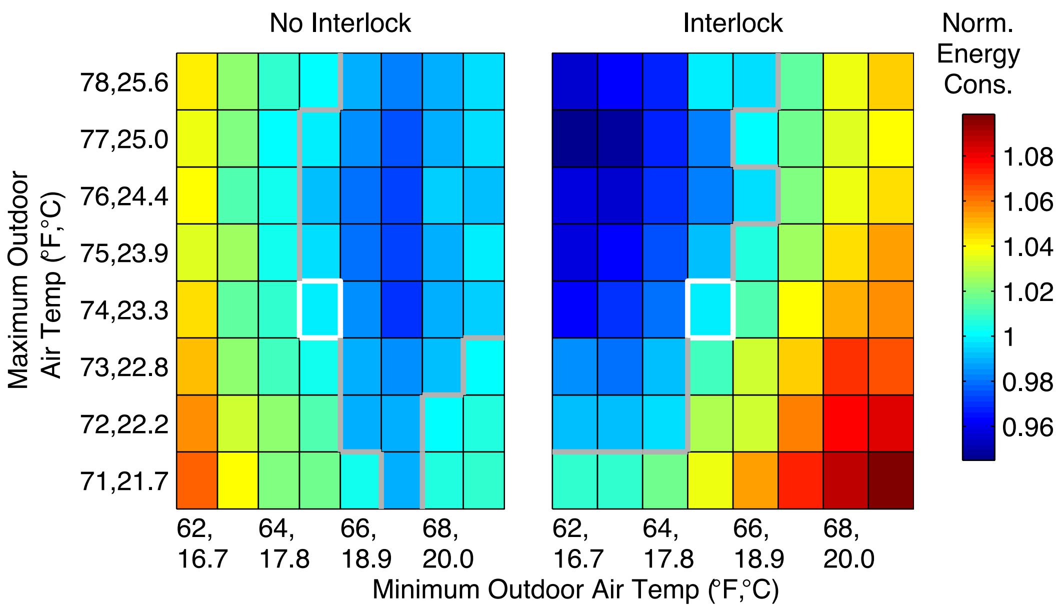

Figure 7.5: Annual energy consumption for no-interlock and interlock controls. Energy consumption for the default case is outlined in white; the light grey line is drawn between regions that indicate savings (value<1), and those that indicate costs (value>1).

comparison from month to month; the values for January are all divided by the default energy consumption for the month of January, values for February are divided by the default consumption in February, and so on. This normalization makes it easier to determine which months have large potential for savings (or increased costs) and which months have small potential for savings. The results in Figure 7.6 show that there is considerable potential for savings in the swing season months of April, May and November, but minimal opportunity in the summer and winter months, suggesting that the existing setpoints are at or near their best values for the summer months. In November, the control pair (67°F,76°F) (19.4°C,24.4°C) uses 23% less energy than the default (65°F,74°F) (18.3°C,23.3°C), and in April, the control pair (69°F,78°F) (20.6°C,25.6°C) saves 24% compared to the default. Savings in these cases come from avoiding window openings during cooler periods and allowing it during warmer periods, thus avoiding pumping energy associated with heating loads; additional savings are realized through reduced heating energy, but are not included in this analysis, as stated above. Note that in Table 7.4, each of the control parameter values considered emerges as the best for one month or another, and there is clearly no single combination of minimum and maximum outdoor air temperatures that leads to consistent energy savings from month to month. This finding suggests that a monthly reset schedule may lead to more savings than a single set of control parameters that is employed year-round. Figure 7.7 shows savings in each month that could be achieved by employing the combination of control parameters that minimizes energy consumption for that month; in each case, the monthly minimum is compared to the monthly no-interlock default value. Figure 7.7 provides a summary of the absolute savings potential in each month as well as a comparison between concurrent and changeover controls. In Figure 7.7 we see that with the one exception of the month of January, a system that incorporates HVAC-interlock has greater potential for savings than a system without. Data presented in Figures 7 through Figure 7.7 is summarized in Tables 4, 5, and Table 7.6. In comparing the monthly data in Table 7.4 with that in Table 7.5, note that the inclusion of an interlock system changes the opportunity for savings in each month; more savings opportunity exists in March, May, September, and October when there is an interlock system, compared to more savings opportunity in April and November when there is not an interlock system. In Table 7.6 the annual performance of a system with the single set of control parameters that minimizes HVAC electric consumption is compared to the performance of a system that uses a reset schedule of the best monthly parameters. Note that in all cases, the addition of an interlock system significantly increases the energy savings potential, and that the simplest case of modifying the existing system by changing two parameters in the control logic leads to 3.1% savings, while the most complex case of a system retrofitted with window-HVAC interlock logic and a reset schedule of monthly control parameters leads to 15.6% savings. This latter case indicates the most involved recommendation that can come out of this study: change the annual setpoints to monthly setpoints, and change the controls of the HVAC and window systems such that interlock behavior occurs, while the former case is as simple as changing two numbers in a direct digital control (DDC) system.

Table 7.4: Monthly summary statistics for natural ventilation study without interlock.

Month	$T_{OA,Min}$ °F, °C	$T_{OA,Max}$ °F, °C	Default-Case Energy Consumption (kWh)	Best-Case Energy Consumption (kWh)	Absolute Savings (kWh)	Percentage Savings
January	62, 16.7	74, 23.3	191	189	2	1%
February	65, 18.3	71, 21.7	147	142	5	4%
March	66, 18.9	78, 25.6	306	285	22	7%
April	69, 20.6	78, 25.6	333	253	80	24%
May	67, 19.4	76, 24.4	404	355	49	12%
June	64, 17.8	77, 25.0	747	723	24	3%
July	62, 16.7	72, 22.2	961	942	19	2%
August	63, 17.2	74, 23.3	1205	1142	63	5%
September	62, 16.7	75, 23.9	381	343	38	10%
October	66, 18.9	73, 22.8	213	199	14	7%
November	67, 19.4	76, 24.4	318	245	73	23%
December	68, 20.0	78, 25.6	142	137	5	4%

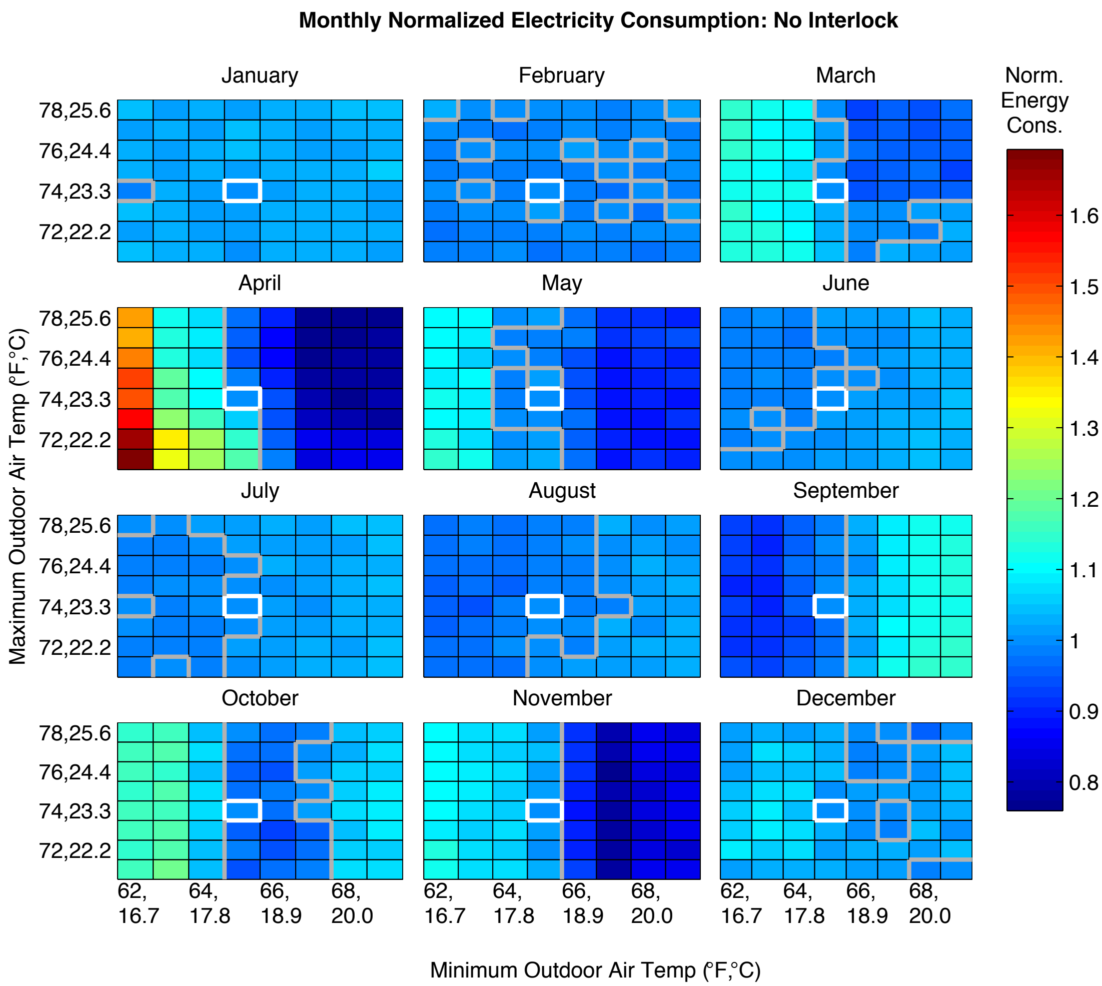

Figure 7.6: Energy consumption for 12 months. Energy consumption for the default case is outlined in white; the light grey line is drawn between regions that indicate savings (value<1), and those that indicate costs (value>1).

Table 7.5: Monthly summary statistics for natural ventilation study with interlock.

Month	$T_{OA,Min}$ °F, °C	$T_{OA,Max}$ °F, °C	Default-Case Energy Consumption (kWh)	Best-Case Energy Consumption (kWh)	Absolute Savings (kWh)	Percentage Savings
January	62, 16.7	74, 23.3	196	191	5	3%
February	62, 16.7	72, 22.2	145	142	3	2%
March	62, 16.7	78, 25.6	288	242	46	16%
April	63, 17.2	76, 24.4	259	238	21	8%
May	64, 17.8	77, 25.0	312	279	33	11%
June	63, 17.2	77, 25.0	695	658	38	5%
July	63, 17.2	76, 24.4	926	901	25	3%
August	62, 16.7	77, 25.0	1149	1072	77	7%
September	62, 16.7	78, 25.6	336	290	46	14%
October	62, 16.7	76, 24.4	186	140	46	25%
November	64, 17.8	72, 22.2	246	225	22	9%
December	65, 18.3	76, 24.4	138	136	2	1%

Table 7.6: Annual summary statistics for natural ventilation study and interlock vs. no interlock comparison.

Description		$T_{OA,Min}$ °F, °C	$T_{OA,Max}$ °F, °C	Default-Case Energy Consumption (kWh)	Best-Case Energy Consumption (kWh)	Absolute Savings (kWh)	Percentage Savings
No Interlock	Single Set of Best Control Parameters	67, 19.4	74, 23.3	5349	5184	166	3.10%
	Monthly Reset of Best Control Parameters	-	-	5349	4955	395	7.40%
Interlock	Single Set of Best Control Parameters	62, 16.7	77, 25.0	4877	4608	269	5.50%
	Monthly Reset of Best Control Parameters	-	-	4877	4513	364	7.50%
Total Potential Savings by Adding Interlock	Single Set of Best Control Parameters	62, 16.7	77, 25.0	5349	4608	741	13.90%
	Monthly Reset of Best Control Parameters	-	-	5349	4513	836	15.6%

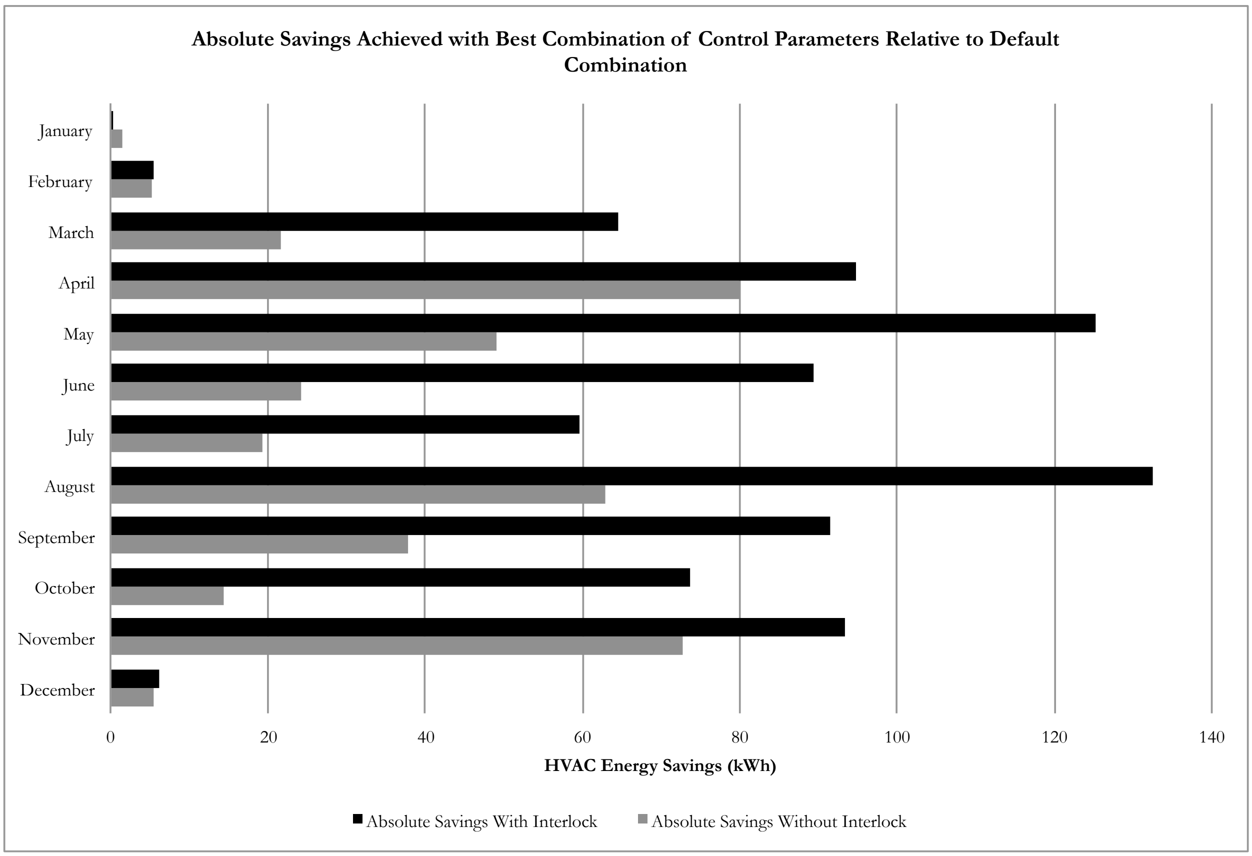

Figure 7.7: Absolute savings achieved with best combination of control parameters relative to default combination. Note that there is no potential savings in January with an interlock system present.

7.4 Conclusion

The automatic window control logic for the RSF building located in Golden, CO was analyzed by means of a parametric study in which sets of values of key control parameters were used in simulations of a calibrated building energy model. The results of each set of simulations gives insight into the energy savings potential of each combination of control parameters, helping to point control designers to a set of control parameters that will minimize energy consumption.

While this study focuses on automatic window control, the methods are applicable to any control strategy with key setpoints and threshold values. Noting that building heating, cooling, and ventilation can be heavily influenced by occupant use of operable windows, a stochastic model of window use was implemented in the building energy model, and the average effect of occupant window use was accounted for by simulating each model a large number of times, and using the average HVAC electric consumption to measure performance. This method allows us to compare the benefits of each combination of control parameters while accounting for the interactive effects between the building systems and occupants. Occupant comfort was accounted for by comparing a weighted sum of comfort conditions computed for each control scenario, and results demonstrated that for this building model, the buildings cooling and heating system consistently maintained comfort conditions regardless of automatic window controls.

Results for the RSF building show that for night ventilation control, opening automatic windows prior to midnight provides little benefit, and that the minimum outdoor air temperature setpoint that minimizes energy consumption is 65°F (18.3°C). For a 24-hour natural ventilation strategy, results showed that the best combination of minimum and maximum outdoor air temperature setpoints that govern automatic window openings is highly dependent on season, and on the presence or absence of a HVAC-window interlock system. In nearly all cases, the addition of an interlock system leads to increased energy savings, and the savings potential from optimizing setpoints is greater when an interlock system is present, largely due to decreased demand on the mechanical ventilation system.

The best-case annual minimum and maximum temperature setpoints are (67°F, 74°F) (19.4°C, 23.3°C) for systems with interlock and (62°F, 77°F) (16.7°C, 25.0°C) for systems without interlock, and lead to annual savings of 5.5% and 3.1%, respectively. If, instead of a single annual combination of control parameters, the controls are augmented by a monthly reset schedule of temperature setpoints, savings potential for a system with interlock increases to a theoretical 7.5%, and savings for a system without interlock increase to 7.4%. The combination of adding an interlock system to a building that lacks one, as well as a monthly reset schedule of window control setpoints leads to a potential 15.6% savings for the RSF building.

Chapter 8

Top Down Results: SMPC and DMPC

8.1 Introduction

A methodology has been developed for optimizing building supervisory control strategies, employing building models that incorporate stochastic models of occupant behavior and serve as the objective function evaluator in a stochastic model predictive control (SMPC) architecture. The SMPC architecture accounts for variability in building performance due to occupant behavior and is shown to generate a sequence of window opening decisions for a mixed mode (MM) building which lead to more robust building performance in the face of occupant window use than a heuristic controller. A set of receding optimization time horizons are described which enable the use of complex building models in simulated SMPC. Results of a case study show that deterministic optimization predicts a 50% increase in building performance, while stochastic stochastic optimization leads to a more conservative and more reliable 33% performance improvement, which takes into consideration the impact of occupant behavior.

8.2 Case Study

The results shown in this section are part of an ongoing natural ventilation controls optimization at the Research Support Facility (RSF) building on the National Renewable Energy Laboratory (NREL) campus in Golden, Colorado, USA. To this end, an energy model of the RSF building has been developed in EnergyPlus, as well as a sub-model which was developed specifically for optimization of the building's automatic window controls. See Figure 5.3 for an aerial photo of the building, and Figure 5.5 for visuals of the energy models. The majority of the RSF building is comprised of office space in the long, narrow, east-west (left-right, in the images) oriented wings, though conference rooms, dining areas, a gym, a library, and a large data center also exist in the core of the building, which is the north-south oriented spine of the building, connecting the office wings to each other.

The simulation study results are presented for the month of October, which is characterized by cooler temperatures and large diurnal swings; night time lows typically range between 0°C and 10°C, and daytime hits range between 15°C and 25°C. During summer months, larger cooling loads are prevalent, but in the month of October we do not expect to see much cooling energy consumption. The results show that the default window controls miss several opportunities to take advantage of natural ventilation for cooling, while the optimal controls exploit these opportunities and prevent the radiant system in the building from switching between heating and cooling several times during the month.

At the time of this simulation study, the automatic windows in the RSF were controlled based on wind speed, outdoor relative humidity, and outdoor dry bulb temperature. The wind speed computations are based on granular checks for gusts over 7.5 m/s, and windows are allowed to open whenever the wind is generally calm, but since weather data was only available in hourly increments, we simply disable window openings whenever wind speed is greater than 7.5 m/s. The relative humidity and temperature based control logic is straightforward; automatic windows are allowed to open whenever outdoor temperature is between 20°C and 23.3°C, and outdoor relative humidity is above 50%.

For each of the manually operable windows, it is assumed that there is one occupant responsible for opening and closing it. This is true in the actual building on the North façade, where there are single-occupancy semi-enclosed offices, and partially true on the southern façade, where there is a desk near each window, but where any occupant is able to open or close any window. In the simulation, each of the 10 occupants that control windows is given a unique occupancy profile generated by the Page occupancy algorithm before the EnergyPlus simulation, and actions are predicted for each occupant for each hour based on the Haldi window-use algorithm. The particular occupancy profiles, parameters of mobility, and long-term absence inputs to the Page occupancy model are drawn at random from the 20 individual sets generated in the model's development ([75]). The probability profiles follow typical 35-hour work weeks - high in the morning and afternoon on weekdays, low during the lunch hour and on weekends, and the mobility parameters are typically 0.12.

8.2.1 Optimization Settings

For the results presented here, a maximum Monte Carlo number (100) was chosen to constrain runtime, though use of the Monte Carlo convergence criteria with $B_L = 50$ and $B_W = .02$ often leads to convergence in less than 100 simulations, and typical Monte Carlo numbers used in similar studies are in the range of 50 to 100, as in [41, 61]. The MPC evaluation was conducted for the month of October, using TMY3 weather data for Golden, Colorado, USA. Execution, planning, and termination horizons are all 24 hours; this results in a 48-hour cost horizon. The decision space is discretized into twelve two-hour blocks per day, always from 18:00 on one day to 18:00 on the following day, which corresponds approximately to the end of the occupied period. In each two-hour block, the optimizer can choose a single decision (whether to open or close automatic windows), resulting in $2^{12} = 4096$ possible window control vectors for any 24-hour period. The optimization algorithm is a modified version of the meta-heuristic particle swarm optimizer (PSO) described in [55]. The SMPC run required 88 hours to run on twelve 2.8 GHz Intel Westmere processors. The goal of the optimization was to minimize cooling energy (electricity consumption for fans, pumps, and chilled water production) while preserving or improving indoor comfort conditions, as in the cost function in Equations 4.5 and 4.6. Before presenting results, the reader should note that the RSF building was designed to operate at net-zero energy conditions, and that it is already near-optimal in terms of minimizing cooling energy. Heating energy was not added to the objective function because at the time of the simulation study, minimizing electrical energy was the primary concern of the building operations staff, and the cost of operating the campus heating system was much lower than the cost of electrical energy for cooling or running HVAC equipment.

8.2.2 Results

Figure 8.1 shows the cumulative energy consumption for fans, pumps, and cooling for the month of October for two cases, a deterministic simulation with default controls, and a deterministic simulation with optimal controls that were selected in a deterministic MPC run. Occupant behavior was neglected for deterministic MPC, meaning that no occupant window-openings occurred in the simulations. Figure 8.2 shows the analogous sets of results for stochastic simulations with default controls and optimal controls as determined by an SMPC run. The first, and possibly most significant contribution of this work is to present results as in Figure 8.2, where we can see the distribution of impacts on building performance from occupant behavior, given default controls and optimal controls. In the box whisker plots, the boxes encompass data between the 25th and 75th quartiles, with the median shown by the horizontal line near the middle of each box; whiskers extend to the last value that falls within 1.5 times the inter-quartile (75th-25th) range; any points outside of this are outliers, and plotted as points. In the deterministic results in Figure 8.1, one could assume some percentage error around the given results, but the stochastic figure shows that the results are not normally distributed, and instead have a long upper tail. In both Figures 8.1 and 8.2, the savings that accumulate when optimal controls are used are characterized by large jumps, for the deterministic results (Figure 8.1) this is obvious on the 18th, 19th, 23rd, and 24th days of the month. In the stochastic optimization results, we see large changes on days 10, 13, 16, 18, 19, 23, and 24 for the default case, and similar but smaller changes on the same days for the optimal case. These large spikes in energy consumption in the simulations with default controls are due to the radiant system changing from heating mode to cooling

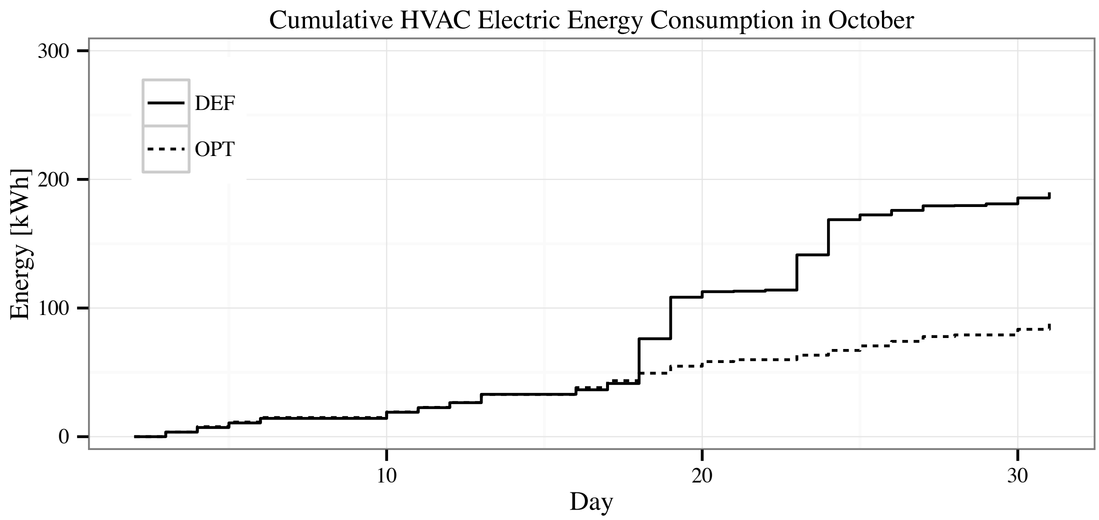

Figure 8.1: Deterministic MPC results: cumulative energy consumption for one month.

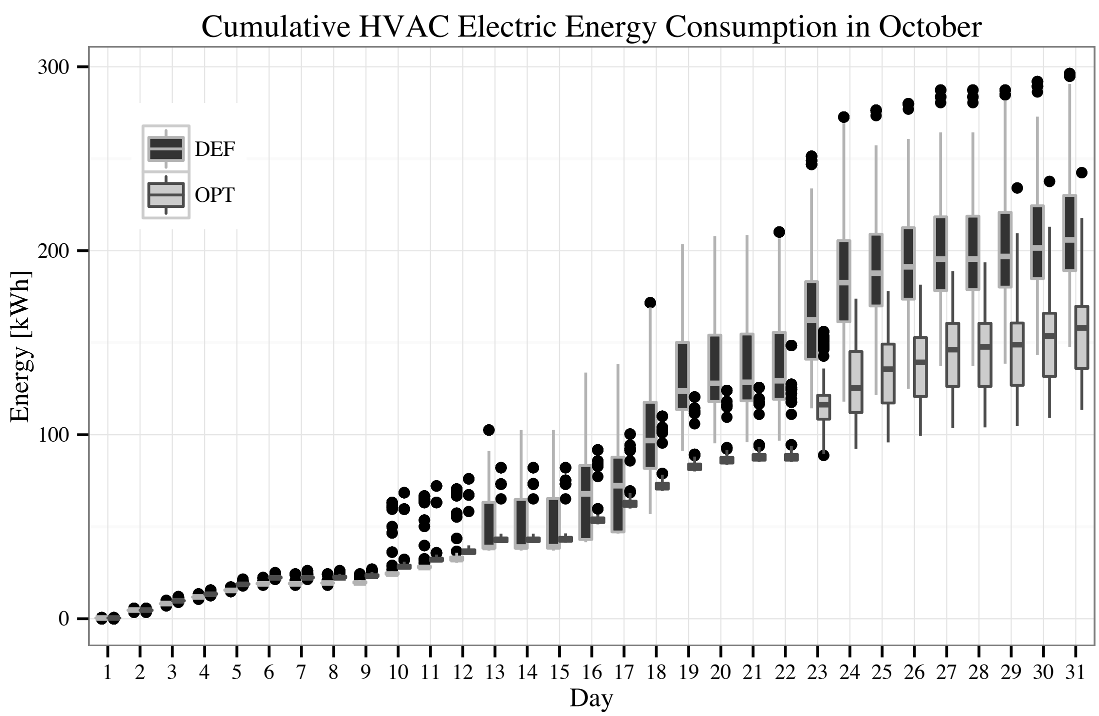

Figure 8.2: Stochastic MPC results: cumulative energy consumption for one month.

mode; with optimal controls in place, windows are opened which lets in cool outside air, prevents the indoor conditions from becoming too warm, and prevents the radiant system from switching between heating and cooling. For the stochastic results, note that the median values of energy consumption do not experience any large jumps until day 16, though the range of potential energy consumption is growing due to occupant behavior.

On other days, there is virtually no difference in performance between the optimal and default controls, because there is no opportunity for natural ventilation to significantly impact the performance of the radiant system. There might be opportunities for fan energy savings if the building's ventilation system were programmed to respond to natural ventilation (i.e. reduce ventilation rates when windows are open), however this is not the case - so savings from fan energy are nonexistent.

The impact of occupant behavior is also very small for the first ten days of the simulation whether default or optimal controls are used, this is due to relatively moderate outdoor conditions, where energy consumption is low in general, however we note that it is possible that the random streams simply lead to lower than typical occupant behavior during this period. The impacts of occupant behavior begin for both the default and optimally controlled buildings on the 10th day in October, however the impacts are slightly different due to the different automatic controls, which lead to different indoor conditions for the rest of the month. An important difference between the stochastic and deterministic results is apparent on day 16, when there is a larger jump in energy consumption for stochastic results (for both optimal and default controls) - this is due to the aggregate impact of occupant behavior over the last two weeks in simulation, and not due to any isolated event.

Figure 8.3 is a histogram of the cumulative energy consumption for the entire month of October for two simulation cases: the building operated with default controls and with optimal controls. Four hundred simulations were conducted with both control sequences (optimal and default) after the MPC run to fill out the distribution of energy for each case. We note that the results of simulations with optimal controls are more tightly clustered, while the results of simulations with default controls are more spread out, indicating that the optimal controls lead to more consistent building performance in the presence of occupant behavior. With default controls, the deterministic result falls near the center of the range of stochastic results, indicating that occupant behavior can increase or decrease building performance in general. For the simulations with optimal controls, at first glance it appears that occupant behavior has shifted energy consumption up, and never shifts it down, but this is not true, since the optimal control sequences found using deterministic and stochastic MPC are slightly different. Some disagreement in results is expected, however the fact that the deterministic result is not even within the range of the stochastic results tells us how important it is to consider occupant behavior during the optimization. The deterministic optimal result may not be valid at all, given that it does not account for the impacts of occupant behavior, while the stochastic result gives a more realistic estimate of potential performance improvements. In both the default and optimal stochastic results, the total energy consumption for the month of October differs by slightly less than a factor of two, which is consistent with the results presented in [41], where the same occupant window-use model is coupled to a single-zone building model, and heating and cooling demands vary by a factor of two. Note that here we do not employ a model of occupant shading device-use since there are no manual shades in the RSF building, however there was occupant shading device-use in the related study; we also do not consider diversity within the occupant pool (i.e. active and passive occupants), while the aforementioned study provides further insight on differences within the occupant pool. In [78], a simulation study that uses a whole building energy model and models of occupant behavior for shading devices and internal lighting shows that the range of lighting energy consumption varies by a factor of two, while the standard deviations of heating and cooling energy use are 9% and 10% respectively, which is lower than the range for this study and other studies that use a single-zone building model, and which incorporate occupant use of manual windows.

Figure 8.4 provides a summary of the savings for the stochastic simulations; the range

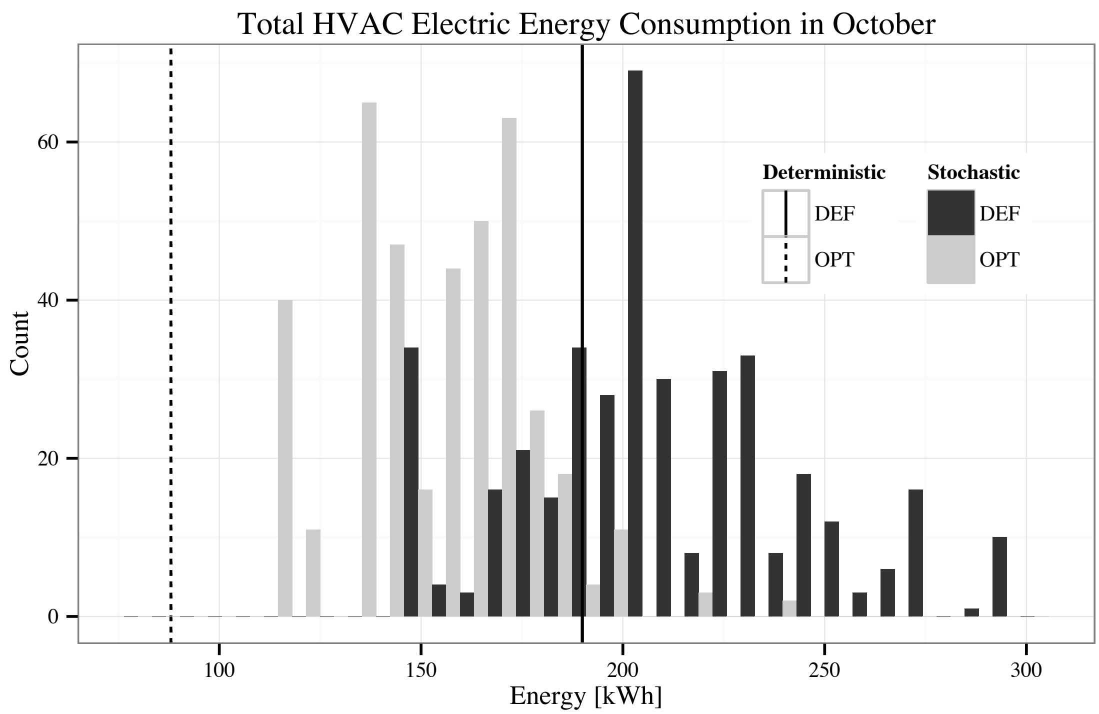

Figure 8.3: Stochastic and Deterministic MPC results: total energy consumption for one month.

of savings is between -68 kWh and 138 kWh, while the average is 55 kWh, and clearly the negative savings estimates are outliers. For the deterministic simulations, energy consumption with default controls was 190 kWh, and with optimal controls was 88 kWh, resulting in a 102 kWh savings prediction. We can deduce from Figures 8.1 through 8.4 that occupant behavior in the RSF tends to decrease the savings potential from controls improvements with MPC.

Figure 8.5 shows the daily energy consumption throughout the month of October. Due to occupant control of manual windows only, the energy consumption varies on each day. On most days, the variation is minimal, indicating that the indoor and outdoor conditions are not conducive to opening windows. On other days however, the variation is quite large, indicating that occupants are opening windows at times when outside conditions are much cooler than the indoor conditions (since the simulation is for October, temperatures are typically cooler outdoors than inside).

Figure 8.6 shows the hourly energy consumption for October 16. Note that here, as in the daily energy consumption displayed in Figure 8.5, there are only a few isolated energy consuming events that account for the majority of the total consumption. On this day, conditions lead to a large energy consumption event at 10:00am that causes the radiant cooling system to turn on.

Figure 8.7 shows the hourly predicted mean vote (PMV) comfort index for October 16: Note that during unoccupied hours the range of PMV indices when optimal BAS window controls are used is allowed to drop well outside of the nominal [-0.5,0.5] range, and that during occupied hours, the PMV drops as low as -0.5, but never lower. The low PMV scores during hours 8 and 9 are a result of BAS windows being opened to prevent an impending cooling event. The presence of the cooling event is obvious in this figure if one looks at the box plots of PMV in hours 10 and 11 for the DEF case; the average PMV value drops dramatically from hour 10 to hour 11, indicating that cooling has occurred. During the simulation, air speed and clothing assumptions that factor into the PMV calculations are

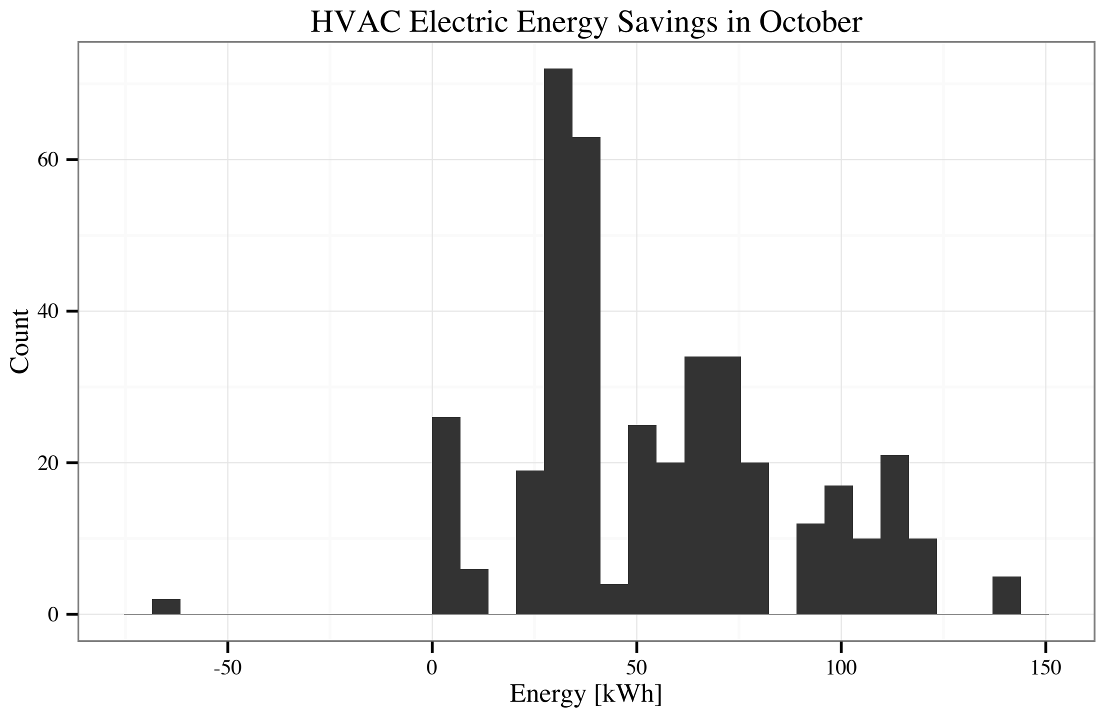

Figure 8.4: Stochastic MPC result: energy savings.

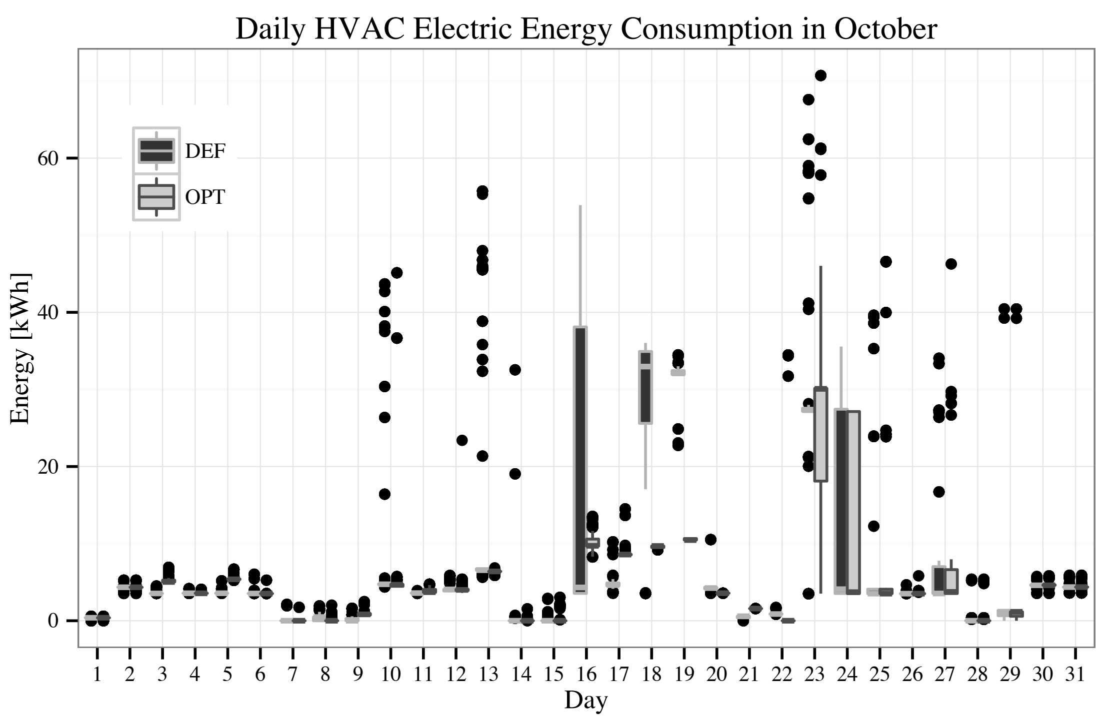

Figure 8.5: Stochastic MPC result: daily energy consumption.

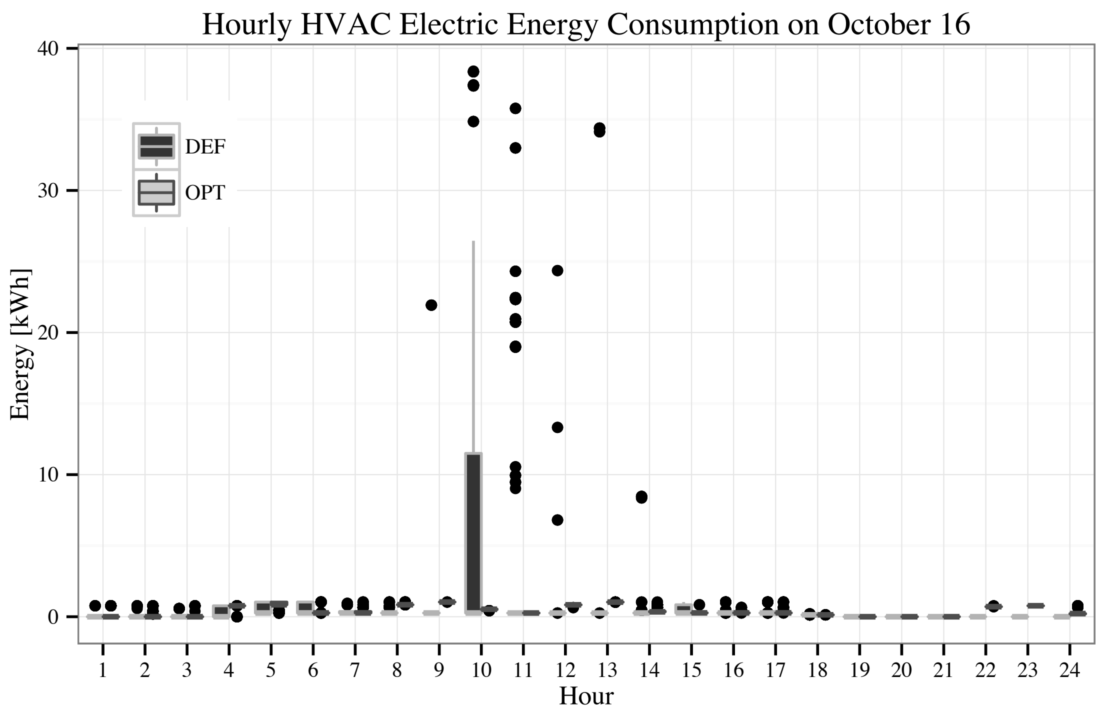

Figure 8.6: Stochastic MPC result: hourly energy consumption.

set to 0.2 m/s, and 0.5 clo during the summer, (May 15-October 15) or 1.0 clo during the winter (Oct 16 - May 14).

Figure 8.8 shows the distribution of energy consumption for a single hour, from 9:00 to 10:00 on October 16. With default BAS window controls in place, the range of potential energy consumption is considerably higher. This is a result of space temperatures rising due to internal loads and solar gains, reaching the cooling setpoint, and causing the cooling system to turn on. With optimal BAS window controls at the helm, windows are opened for two hours before 10:00, which causes indoor space temperatures to drop significantly, preventing the cooling system from turning on. Note in Figure 8.7 that the indoor comfort index drops to -0.5 during occupied hour 8. In this case, the window controls are acting like an outside air economizer with a bit of predictive power. The DOAS in the building model is only intended to provide fresh air, so while it could potentially perform the same function that the windows are, it is not designed to. Instead, cooling and heating come from the concrete slab in the ceiling, and in order for the radiant system to respond to this cooling event, it has to completely switch from providing hot water to a typically warm slab, to providing cool water to cool the slab and then the space. It is this switch that draws such a large amount of energy in the higher cases in Figure 8.8.

In Figure 8.8, we see that the DEF case can experience large energy consumption, depending on the behavior of occupants, while the OPT case has successfully avoided any large energy consumption event. Note in the subplot in Figure 8.8 that the OPT results do still contain some variation, but the variation is limited to a much smaller range than in the DEF case.

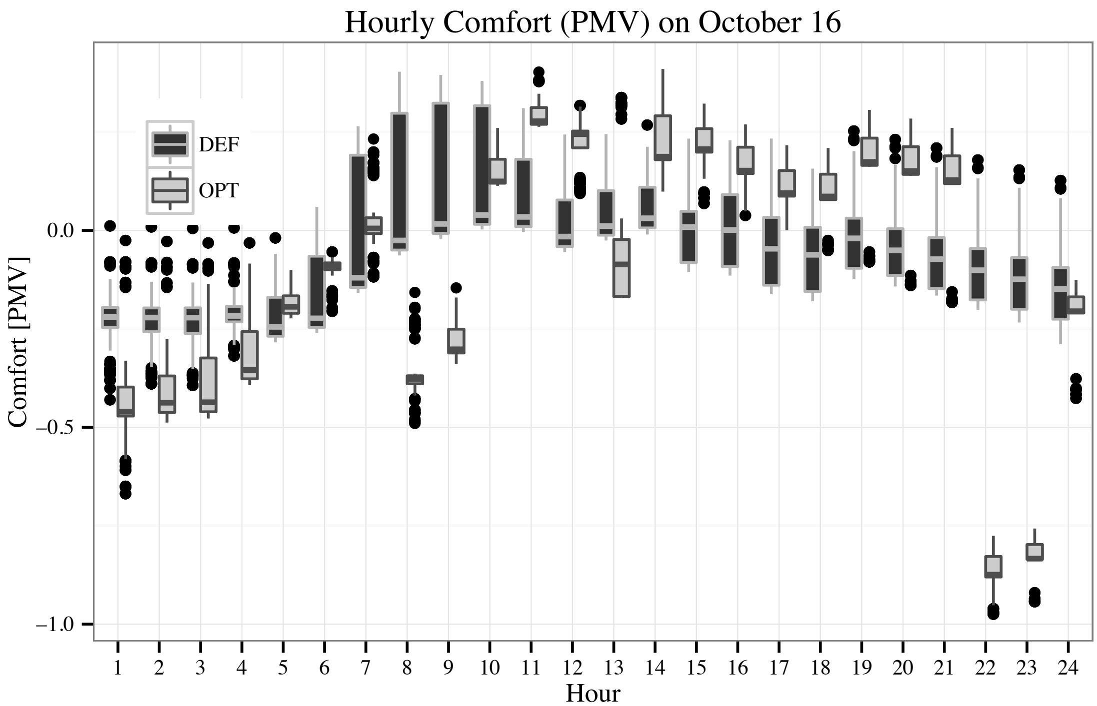

Figure 8.7: Stochastic MPC result: hourly comfort.

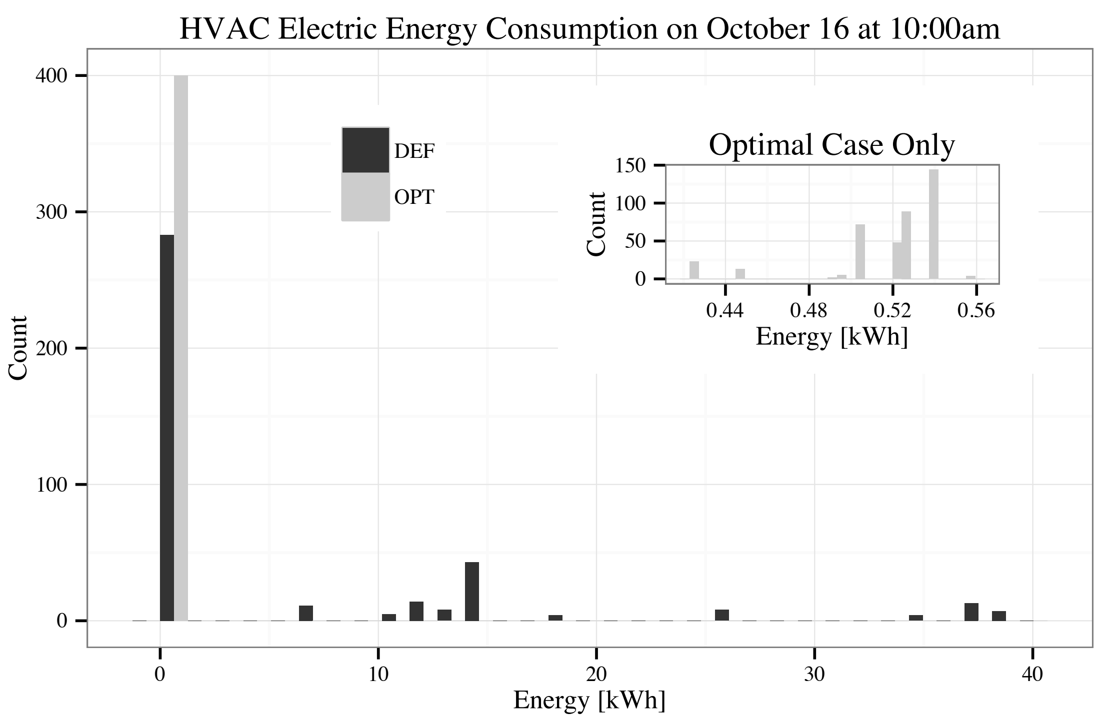

Figure 8.8: Stochastic MPC result: energy consumption for one hour.

8.3 Conclusion

A methodology has been developed for optimizing building supervisory control in simulation, in which building models developed in EnergyPlus can incorporate stochastic models of occupant behavior and serve as the objective function evaluator for a stochastic model predictive controller. The SMPC architecture accounts for different levels of variability in building performance due to occupant behavior, and is shown to provide control setpoints which lead to more consistent building performance in the face of occupant window use. A set of time horizons are described which enable the use of complex building models in simulated SMPC. Results of a case study show that stochastic optimization leads to more conservative performance, but guarantees that performance with some level of confidence, as compared to deterministic optimal control, which does not consider any uncertain disturbances and only provides a point estimate of performance. The optimal data sets that are created by this process provide useful insight into the performance of existing control strategies, and can be used to create better control strategies. In the case study results, the SMPC result highlighted several occasions where an existing heuristic control strategy was not performing as well as it could.

Chapter 9

Rule Extraction

In [26] and [63], offline optimization results are used with different statistical modeling techniques to develop simple control rules that approximate the performance of optimal controllers. In this chapter, CARTs are used with results of multiple deterministic optimizations of window controls in the simplified RSF building model.

In [63], May-Ostendorp shows that control rules trained on a dataset of optimization results for the cooling season do not perform well when tested in a swing season, whereas in [26], Domahidi demonstrates that a rule trained on the results of multiple optimizations in different climates works well when tested in each individual climate. Following from these initial findings, we look a bit deeper into training CARTs on optimal results from different times of year, and from different climates.

First, in Section 9.1, we lay out the structure of several different CART training processes that were used to generate dozens of CARTs for each optimal result and select the best one based on closed-loop validation. Second, we put the methods to work training CARTs on datasets corresponding to each month during the swing seasons and the summer, as well as the aggregate dataset of all seven months, then pit the rules from each month against each other in a contest to see which rule can perform the best in any or all other months, and to tease out the generalizability of CARTs trained on limited datasets. Third and finally in Section 9.4 we explore the robustness of rules with respect to climate, training rules using optimization results from 42 different locations and comparing the performance of each rule in every location.

9.1 Rule Extraction Techniques

The general approach to training CARTs in this work begins with preparing the data for rule extraction. The optimal control signal in this case is always the status of automatic windows in the RSF. In its raw form, this signal takes the form of a binary (0,1) sequence, corresponding to closed (0) or open (1). In both [63] and [26], binary signals are the response variable predicted by CARTs, and performance of the CARTs is shown to be strong in both cases, but later we show that transforming this response from binomial to multinomial can help in closed loop testing.

Given our response variable, the window status, we can collect all of the potential predictor variables that we think may be relevant; some past and some predicted outdoor weather parameters, as well as a selection of current-time indoor environment parameters that were used in the results presented here are listed in Table 9.1.

Armed with a set of predictors and a response, we follow the methodology set out in [42] to first grow a tree to an acceptable level of complexity, then prune it back to a less complex structure using the typical one standard error (1-SE) rule of thumb. An example CART dendrogram is given in Figure 9.1, and the pruning process can be understood by examining the complexity chart in Figure 9.2, where the horizontal line corresponds to the value of relative error that is one standard error larger than the lowest value for any split in the tree. The rule of thumb specifies that the appropriate number of splits is the one with the highest relative error that still falls within one standard error of the minimum. In this case, the solution is trivial and corresponds to the final (10th) split, so the pruned and full tree are the same. In Figure 9.1 this is indicated by the color of the trunk in the dendrogram, where any **split** that is fed by a red branch would be pruned away by the one standard error rule. In this case, every split is fed by a green branch, so every split is retained in the pruned model.

Table 9.1: Weather parameters used in rule extraction.

Time Period	Parameter	Value	Sum	Min	Max	Mean	Swing
Past 24 Hours	Solar Radiation		•				
	Relative Humidity					•	
	Outdoor Dry Bulb			•	•	•	
Future Days (0,1,2)	Outdoor Dry Bulb			•	•		•
Future Hours (1,2,3,4)	Outdoor Dry Bulb	•					
Current Value	Outdoor Dry Bulb	•					
	Outdoor Dew Point	•					
	Solar Radiation	•					
	Wind Direction	•					
	Wind Speed	•					
	Total Sky Cover	•					
	Opaque Sky Cover	•					
	Precipitation	•					
	Occupancy	•					
	Indoor Air Temperature	•					
	Indoor Relative Humidity	•					
	Indoor Humidity Ratio	•					
	Indoor CO_2 Concentration	•					
	Hour of the day (1-24)	•					
	Weekday (1-7)	•					

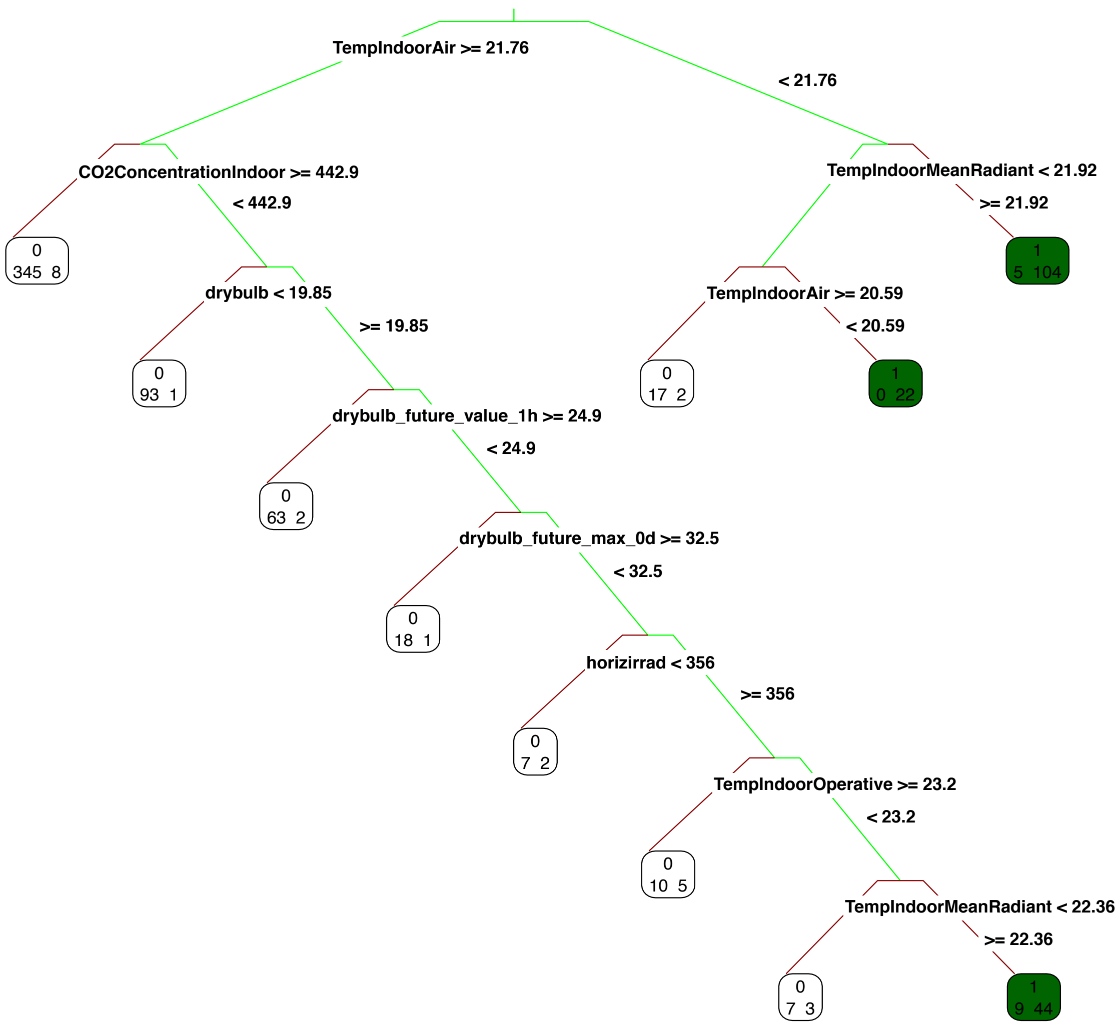

Figure 9.1: Month of August binary CART. White and green boxes are the 'leaves' of the tree, representing closed and open windows, respectively. Numbers in the leaves indicate the response level (0 or 1), as well as the count of each response that falls in each leaf.

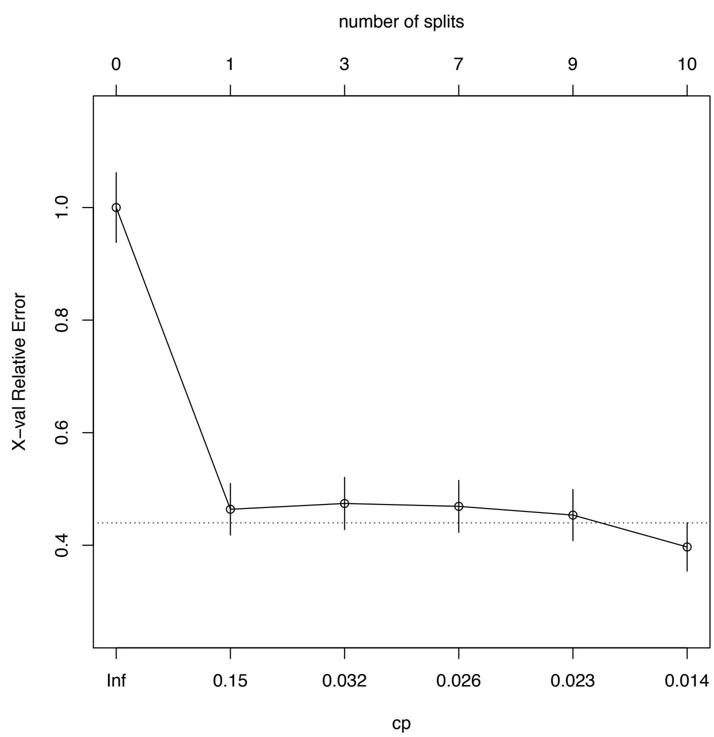

Figure 9.2: Month of August initial binary CART complexity diagram.

At this point, we have a basic strategy for developing basic or what we refer to as **original** CART models in two forms, the fully grown or full tree, and the pruned version. Next we will explore different options for optimizing CARTs, and for massaging the data in our response variable to better approximate real-world building controls.

9.1.1 Evaluating CART performance

In the next section a method of optimizing CARTs is discussed, here we describe the specific metric used to score CART models using their own training data and when tested in EnergyPlus simulations. The ranked probability skill score (RPSS) is a common metric for rating the performance of categorical prediction models. The RPSS compares the predictive ability of a selected model (the ranked probability score, or $\overline{RPS}$) to the predictive ability of an unskilled model $\overline{RPS}_{unskilled}$, and can take values between $-\inf$ and 0. The RPSS is computed in Equation 9.1

$$RPSS = \frac{\overline{RPS} - \overline{RPS}_{unskilled}}{0 - \overline{RPS}_{unskilled}} \tag{9.1}$$

where $\overline{RPS}$ is given by

$$RPS = \frac{1}{M-1}\sum_{m=1}^{M}\left[\left(\sum_{k=1}^{m} p_k\right) - \left(\sum_{k=1}^{m} o_k\right)\right]^2 \tag{9.2}$$

in which p_k is the predicted probability in forecast category k, and o_k is an indicator (no=0, yes=1) for the observation in category k, and M is the number of discrete response categories. The unskilled forecast, $\overline{RPS}_{unskilled}$, could be as simple as a proportional probability for each response category.

While RPSS is appropriate for evaluating CART models in open-loop tests, i.e. using the raw training data, it is less appropriate for evaluating the performance of the CART in an EnergyPlus simulation. For this case, closed loop testing, we can simply use the same metrics that were used to evaluate control performance in the MPC optimizations in Chapter 8, and in the parametric studies in Chapter 7: HVAC electric consumption, indoor comfort conditions, or the combination of two in the objective function in Equations 4.5 and 4.6.

9.1.2 Optimizing Priors

From [63], we know that we can improve the performance of a CART by giving preferential treatment to different levels of the response variable, but it is unclear whether we should prioritize incorrect *window* = *open* predictions, or incorrect *window* = *closed* predictions, though May-Ostendorp suggests that missing the right window **opening** decisions is what leads to less than optimal performance for CART-based window control rules.

In an attempt to maximize the performance of a CART model, we set up an optimization routine wherein the objective function evaluator trains a CART model given some prior probabilities for each response class (chosen by the optimizer), computes the RPSS, and returns the RPSS score as the objective function value. By optimizing the priors to maximize the RPSS score, we guide the tree growth away from misclassification, and closer to a perfect prediction - but for its own training data. The results of an RPSS score optimization, further described below, are shown in Table 9.4. The table contains RPSS scores for various extracted rules for a single set of training data, of note is that the RPSS score generally increases as we move from the first column to the third, which contain scores for original and optimized CARTs, respectively.

Bearing in mind that the end-goal of this rule extraction exercise is to arrive at a pragmatic, usable control logic, we would prefer to have pruned models, and experience tells us that full CART models, while still understandable, can be a bit unwieldy and still difficult to implement if they are excessively large. To that end, we have implemented our CART optimization routine in two forms, in the first form we compute the RPSS for fully grown trees, and return this score to the optimizer - so we are essentially training fully grown trees. In the second implementation, we grow a CART, prune it back per the 1-SE rule, and compute the RPSS of the pruned tree; in this case we are optimizing a pre-pruned tree by adjusting the prior probability values.

At the conclusion of a single CART growing season, we are left with six trees to choose from, a full and pruned version of the original CART, and similar versions of an optimized full tree and an optimized pruned tree, all of which are listed in Table 9.4.

9.1.3 Multinomial Responses

Initial forays into the world of CART training with RSF results showed an interesting trend, which is best explained through an example. Consider the RSF building on a perfect day for pre-cooling; the early morning hours are cool, and the daytime will be hot - so the building can benefit from opening its windows prior to occupancy, harvesting some cold outside air and dumping its heat to the outdoors. In an MPC run, the optimizer recognizes this and takes advantage by opening windows during the early morning hours and closing them during the warm afternoon. The result is an optimal dataset that contains several hours where the indoor temperature is cool, and the windows are open.

Now step forward to the rule extraction process when a CART algorithm is faced with choosing what variables to partition the response variable by, and the CART algorithm notices that whenever windows are open, the indoor temperature is cool - and it chooses to recommend **opening** windows when the **indoor** temperature is cool. In an open loop test - where the CART is asked to predict window position given its own training data, it scores well - choosing to open windows when the indoor temperature is low.

When we implement this control rule in a closed-loop simulation, however, the rule performs poorly, because it got things backwards; the CART captured the relationship well - but not the correct direction of cause and effect. In a closed loop simulation, the CART rules open windows whenever it is cool inside, regardless of the outdoor conditions, which leads to increased heating demand in the swing season, and, sadly, increased cooling demand in the summer.

While it is clear to an engineer that the open window caused the cool indoor conditions during a cool outdoor period, the CART algorithm misinterprets the causality. Therefor we describe some methods for changing the response signal from binary to multinomial so that the CART can (hopefully) see things more clearly. Start by considering the difference between an open window, a closed window, an open**ing** window, and a clos**ing** window. In the binary signal, only the state of the variable is captured - not the **change** in state. If we can capture the extra information on the **change** in state of the windows, and provide this information to the CART algorithm, perhaps it will see that when windows are open**ed**, it is warmer inside, and when windows are clos**ing**, it is cooler inside.

Tables 9.2 and 9.3 will help in explaining how the binary response signal was transformed into a multinomial signal for the CART. After the time index in the first column of Table 9.3, we see the original binary response, followed by eight columns of different multinomial transformations, which are intended to capture the **change** in state of the original response.

In the first transformation, Multinomial-1, the change in state is associated with the first hour that the response is in a changed state relative to the previous time step. In the next column, Multinomial-2, the change in state is associated with the time-step **before** a new state is realized. All of the first four columns represent 3-level multinomial responses, where the state of the window is interpreted as opening (-1), closing (+1), or static (0), which could indicate an open window **or** a closed window.

In columns 5-8, we introduce another level of detail, differentiating between open and closed windows, instead of only considering the window to be static. We now have eight different transformations of the original binary response, and we consider all eight because we know that the optimizer opens and closes windows considering present and future conditions - and we want to enable the CART algorithm to mimic the optimizer as closely as possible. Perhaps the optimizer chose to open windows at time index 4 in the binary response in Table 9.3 solely based on conditions during hour 4 - but maybe something during hour 3, i.e. the building heating up, led to the opening decision, and thus the open window during hour 4.

Without knowing a priori whether to use the hour before the change in state, or the hour after the change in state, or whether it should be the hour before for opening decisions,

Table 9.2: Binary and multinomial response case value key.

	Window State		
Value	Binary	Multinomial 1-4	Multinomial 5-8
0	closed	static	closed
1	open	opening	opening
-1		closing	closing
2			open

Table 9.3: The original binomial response, and multinomial variants.

Time Index	Binomial	Multinomial							
		1	2	3	4	5	6	7	8
1	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0
3	0	0	1	0	1	0	1	0	1
4	1	1	0	1	0	1	2	1	2
5	1	0	-1	-1	0	2	-1	-1	2
6	0	-1	0	0	-1	-1	0	0	-1
7	0	0	1	0	1	0	1	0	1
8	1	1	0	1	0	1	2	1	2
9	1	0	0	0	0	2	2	2	2
10	1	0	-1	-1	0	2	-1	-1	2
11	0	-1	0	0	-1	-1	0	0	-1
12	0	0	0	0	0	0	0	0	0

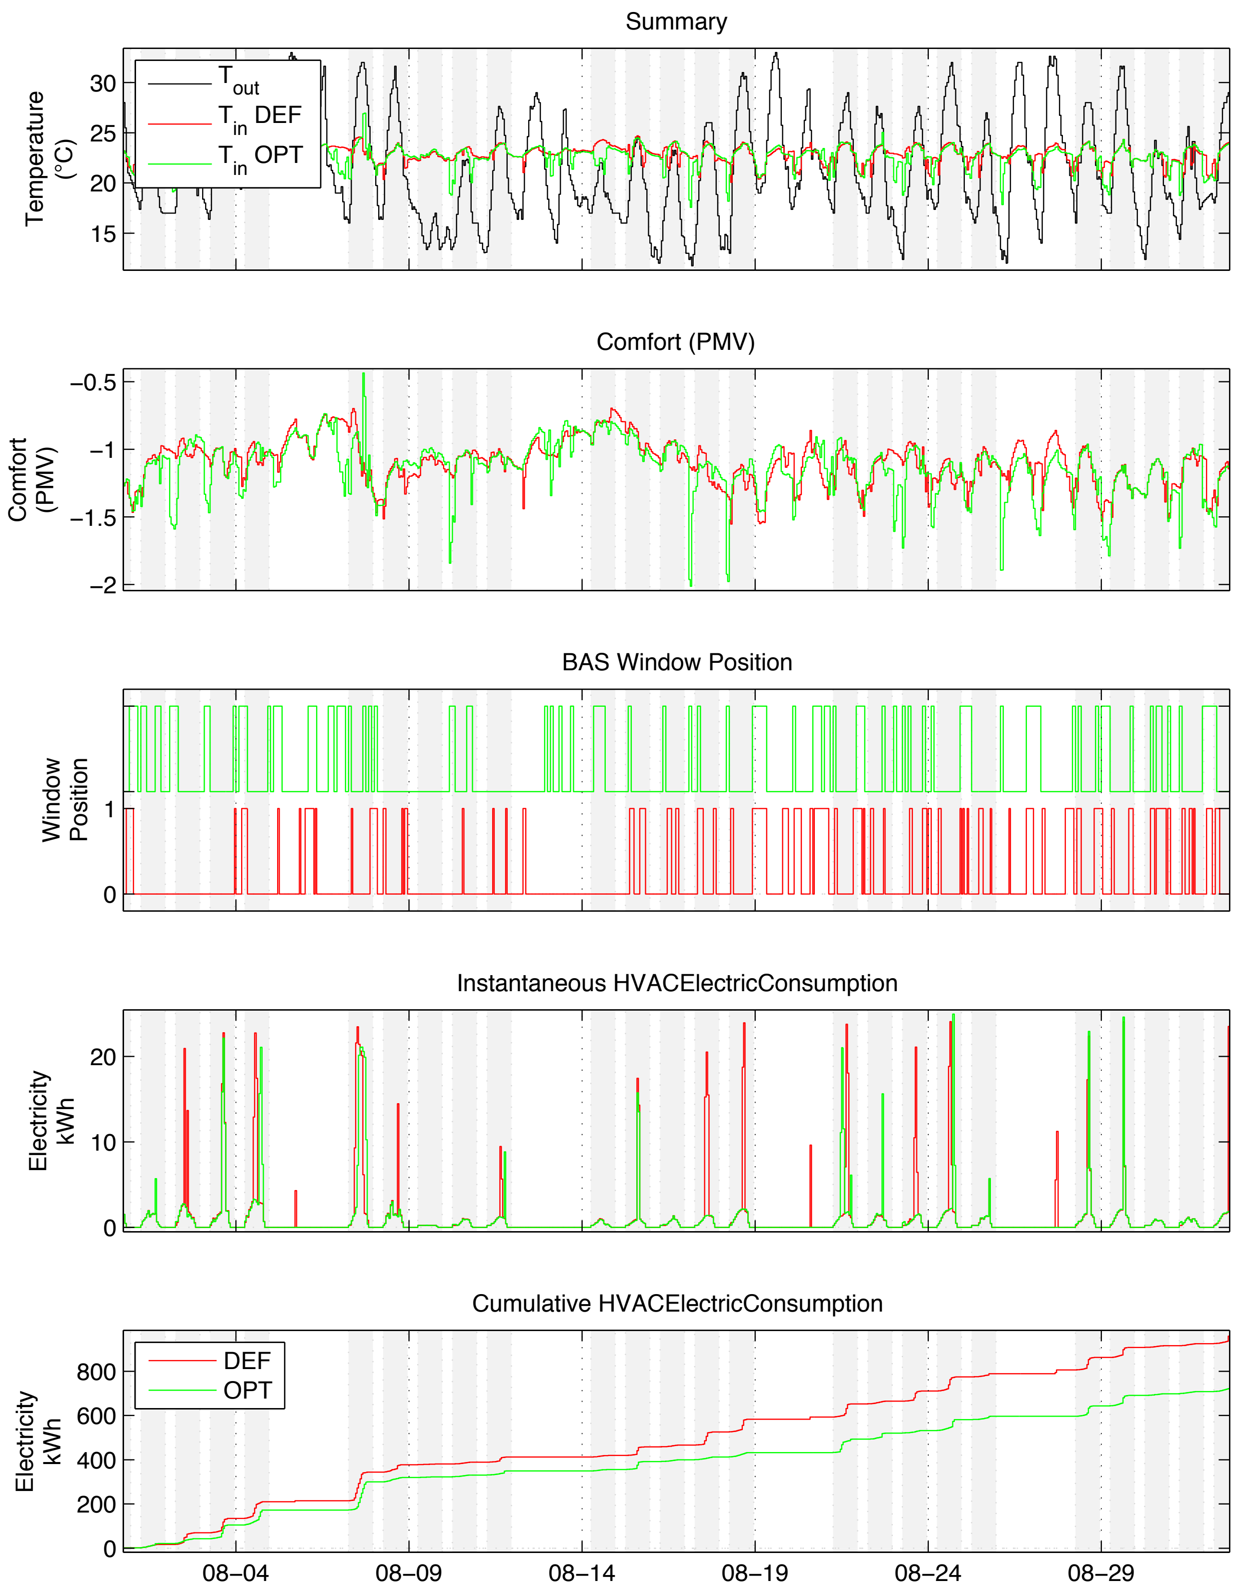

Figure 9.3: Month of August optimization results summary.

and the hour after for closing decisions, or vice-versa, we try all eight permutations as they are laid out in Table 9.3. Recall that we had six different methods for growing and pruning trees in Section 9.1.2, combining those six methods with the nine different response classes, we have 54 different potential options for CART models for a given data set.

9.2 Results for a Single Month

Table 9.4 gives a summary of the RPSS metric for each of 54 different CARTs that were trained on one month of data that came from optimizing the RSF model for the month of August. A summary of the initial MPC optimization results is given in Figure 9.3, where the default and optimal control signals are displayed in the middle panel. When CART derived rules are embedded in the EnergyPlus simulation, most of them perform poorly in spite of some relatively high (0.7-0.8) RPSS scores in open loop testing. Computing the RPSS for the closed-loop simulations results in the scores presented in Table 9.5, all of which are below zero - indicating that they perform worse than an unskilled model. However, as revealed by further analysis, a low open-loop RPSS is not always indicative of poor performance in closed-loop simulations.

Table 9.4: August CART open loop validation RPSS scores.

	Original		Optimized - Full Tree		Optimized - Pruned Tree	
Response	Full	Pruned	Full	Pruned	Full	Pruned
Binary	0.77	0.77	0.84	0.55	0.84	0.82
Multinomial 1	0.34	0.00	0.40	-0.02	0.23	-0.01
Multinomial 2	0.21	0.00	0.34	0.00	0.27	0.00
Multinomial 3	0.28	0.00	0.49	-0.02	0.28	0.19
Multinomial 4	0.17	0.00	0.38	-0.01	0.17	0.00
Multinomial 5	0.63	0.43	0.72	0.63	0.67	0.67
Multinomial 6	0.51	0.33	0.56	0.25	0.46	0.46
Multinomial 7	0.52	0.23	0.58	0.21	0.50	0.15
Multinomial 8	0.65	0.50	0.65	0.57	0.66	0.63

Table 9.5: August CART closed loop validation RPSS scores.

Response	Original		Optimized - Full Tree		Optimized - Pruned Tree	
	Full	Pruned	Full	Pruned	Full	Pruned
Binary	-2.67	-2.67	-1.87	-2.95	-2.06	-2.95
Multinomial 1	-0.45	-3.08	-2.98	-3.08	-0.32	-3.08
Multinomial 2	-3.08	-3.08	-1.79	-3.08	-3.10	-3.08
Multinomial 3	-3.01	-3.08	-1.87	-3.08	-3.08	-0.32
Multinomial 4	-2.78	-3.08	-0.49	-3.08	-2.78	-3.08
Multinomial 5	-3.08	-0.77	-0.55	-3.08	-2.63	-3.08
Multinomial 6	-0.43	-3.08	-2.97	-3.08	-3.08	-3.08
Multinomial 7	-3.08	-3.07	-3.03	-3.07	-2.99	-3.06
Multinomial 8	-3.08	-3.08	-3.08	-3.08	-3.08	-3.08

Because the actual conditions evolve differently in the closed-loop simulation than what was present during the optimal simulation, computing the RPSS for these two cases is not a very good indicator of performance. Instead we can observe the energy and comfort performance of the models in simulation, as in Figure 9.4, where the open loop RPSS and the closed loop objective function are juxtaposed in the bottom panel. Note that the actual objective function value is extremely large due to the weighting coefficients introduced in Section 4.3.4, and have been log-transformed for the chart.

Figure 9.4 provides a number of insights. Note that the $\times$ symbol in the chart corresponds to a 'trivial' controller, or one which leads to closed windows for the duration of the simulation. A number of extracted rules fell into this group, and were combined into the single $\times$ marker. Additionally, the CART model that performed the best in the closed loop simulation is denoted by a black circle. The solid horizontal black line corresponds to the performance of the model with default controls (the controls currently in use at the RSF), and the dashed line corresponds to the performance with the optimal control signal.

First observing the colors of the points in Figure 9.4, we can see some trends in the relative performance of different response cases; CART models trained with the original binary response signal clearly achieve the highest RPSS values on average, with multinomial-5, 6, and 7 also scoring well. Next, looking at the shapes we can see that there is no consistent trend in terms of tree-growth; optimizing the priors to achieve high RPSS scores does not correspond to better performance in closed-loop simulation. One interesting trend however is that pruned models (hollow shapes) often outperform their 'full' counterpart (filled shapes) according to closed-loop metrics, in spite of lower open-loop RPSS values.

We can also see in Figure 9.4 that some extracted rules outperform the optimizer **and** the default controller in terms of the comfort metric (middle panel), but none are able to match the optimal or default controller in HVAC electric consumption (top panel). A final note is that the optimizer performs worse than the default controller in terms of the objective function, which is clearly due to the worse-than-default comfort value seen in the middle panel. Given that our objective function was intended to guide the optimizer towards improving HVAC **and** comfort performance, how is this possible?

The answer to this question comes in two parts. For the first part, recall the discussion of divergent thermal histories and Figure 4.3; at some point in the MPC run, the optimizer chooses a different control sequence than the default controller. From that point onward in the MPC run, the indoor conditions in default simulations and optimal simulations will be slightly different, and at some point the optimizer will be stuck in a situation where it can not achieve comfort conditions that are equivalent to those in the default simulation. At this point, the optimizer still tries to outperform the default controller but can not, and has to choose the least 'worse-than-default' control sequence, or choose the least of all evils. It is under these conditions that the optimizer is able to trade energy consumption for comfort conditions, depending on which is 'cheaper' in terms of the objective function, leading to results like those in Figure 9.4.

The first part of this answer can be attributed to the near-sightedness of the MPC process. While we consider the impact of control decisions over the future 48 hours, we can not consider the infinite impacts, so there are inevitably cases where a decision made early in a simulation can have a slightly negative effect far later in the simulation. Short of providing a multiple-week cost horizon for the optimizer, this is an unavoidable consequence.

The second reason for the occasional failure of the optimizer to produce perfect long-term results is its lack of access to anything besides windows. All of the optimizations have only considered opening and closing automatic windows, and as such the optimizer can not force the ventilation, cooling, or heating systems to actively respond.

While Figure 9.4 gives the final aggregate metrics for the CART performance for the month of August, it does not give any insight into the dynamics that led to the HVAC, comfort, or objective function values. What did the different CART models do to achieve their performance? The initial binary model in Figure 9.1 allowed windows to open whenever the indoor air temperature was less than 21.76°C, and the mean radiant temperature was above 21.92°C (the upper-rightmost green leaf); other window-opening decisions are also based primarily on indoor temperatures, though current and future outdoor temperatures and solar radiation are considered.

The CART that scored the best in terms of the objective function value was the unpruned version of the Multinomial-7 response case, where prior probabilities were optimized to maximize the performance of a pruned model. This unpruned version is more complex, and includes heating setpoints and indoor CO_2 concentration as predictors. The closed-loop response of each of the top 20 CART models is plotted in Figure 9.5; note that there are very few window-openings in general, that the 'trivial' case, of windows being closed the entire time, comprises 9 out of the top 20 performers, and that none of the CART rules seem to come close to re-creating the optimal controller's performance.

While these findings are less than exciting, they serve to illustrate the process which is employed here repeatedly. Given a set of optimal results, 54 different CARTs are created with varying response, growth, and pruning characteristics. Each CART is tested in a closed-loop simulation, and a 'best' rule can be selected based on any of the standard building performance metrics. Next, this process is applied to multiple month-long optimization results, and to an aggregate seven-month dataset.

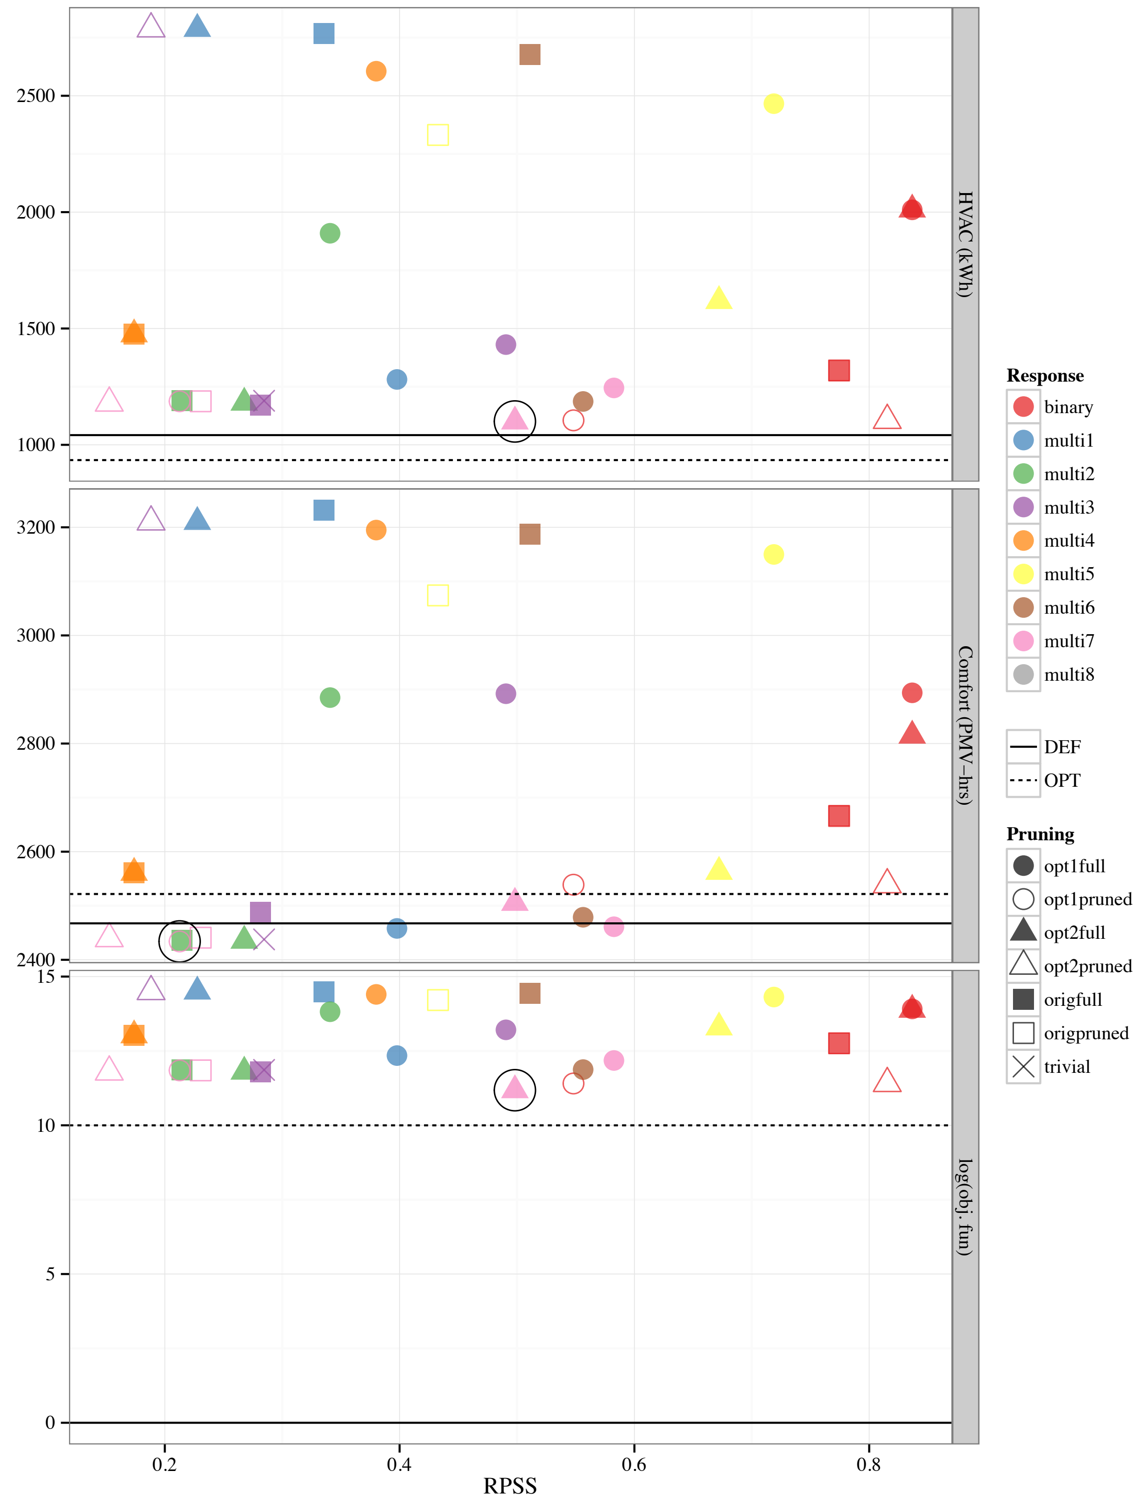

Figure 9.4: Month of August open loop vs. closed loop performance. Black circles denote the 'best' closed loop performance for each metric. The × marker represents a simulation with the windows closed.

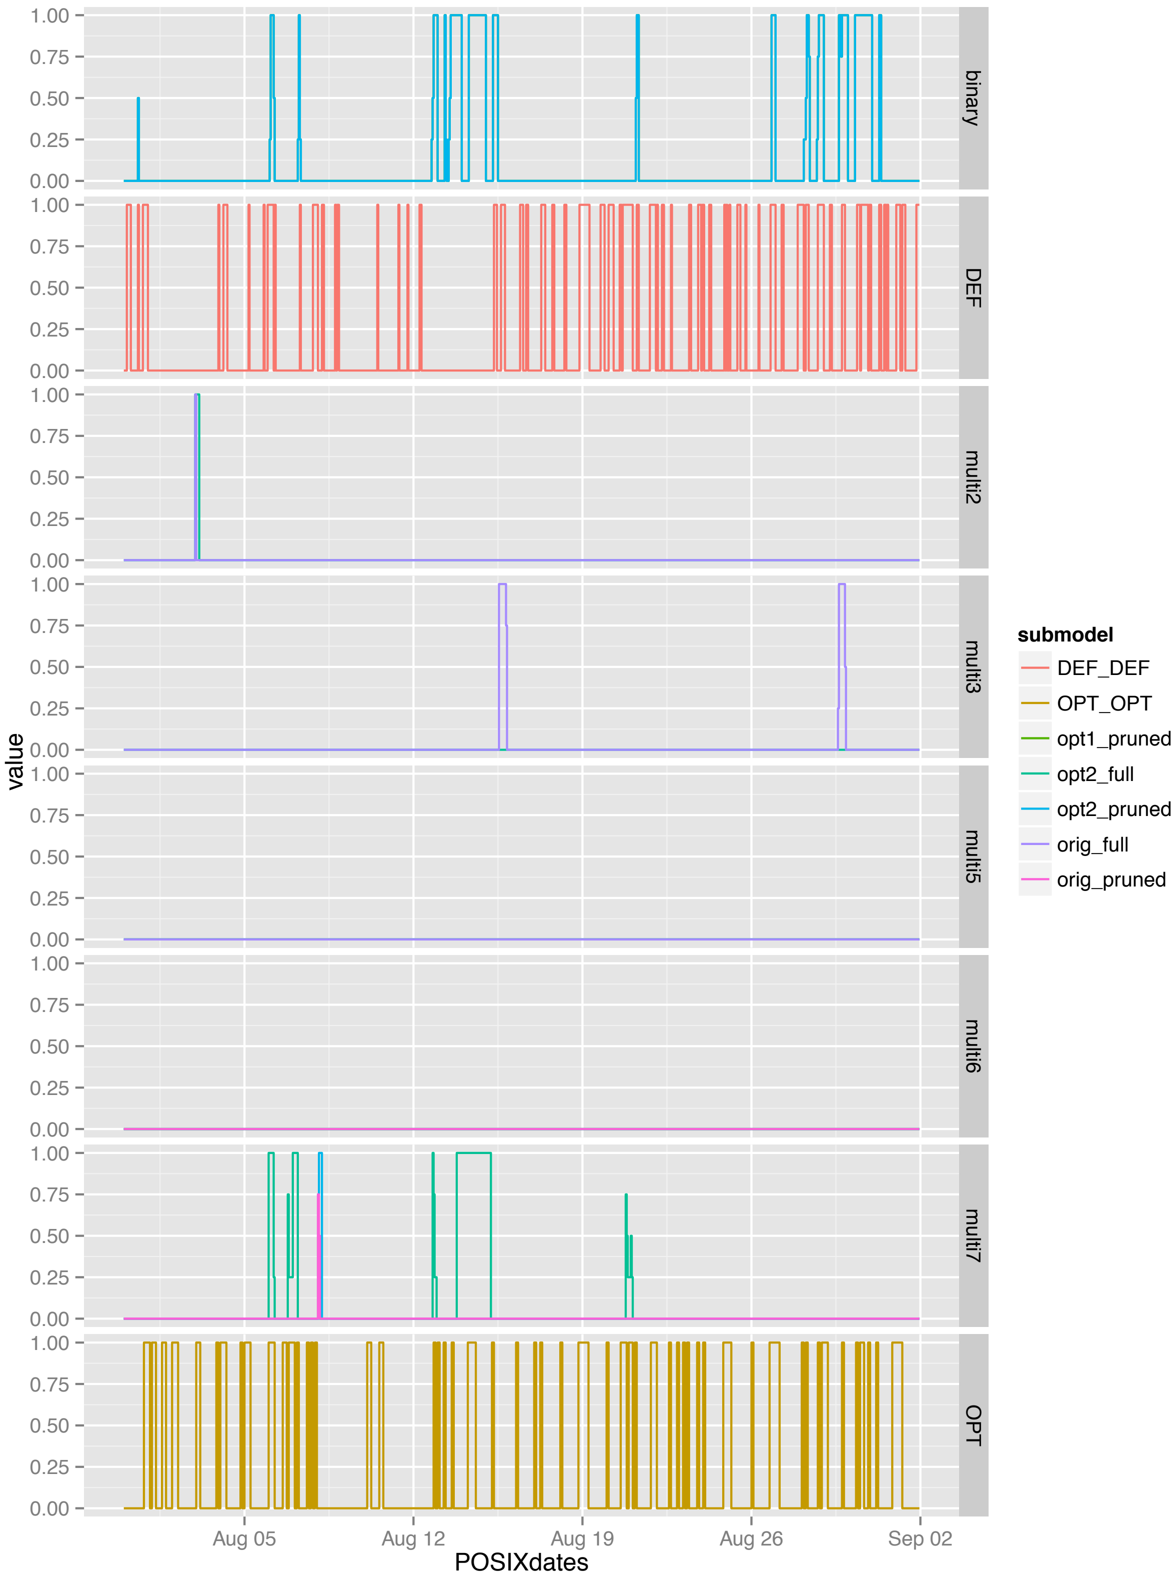

Figure 9.5: Month of August rule extraction closed loop response.

9.3 Rule Extraction for Multiple Months

After looking at the results of optimizing controls for a single month, and applying rule extraction to the limited one-month optimal dataset, we can do the same for multiple months. Recall that for the month of August, the full version of the multinomial-7 response case with priors optimized to maximize the performance of a pruned model was the best performer in terms of the objective function. Figure 9.6 shows the 'best' performing type of CART for each of seven individual one-month datasets, as well as the aggregate seven-month dataset, which includes all months from April through October. While there is no definite overall standout, a few observations are worthwhile. First, no single type scores the best in both comfort performance and HVAC energy consumption, though several score best in both HVAC and the overall objective function, which is an indication that HVAC energy has a larger impact on the overall objective function value. Second, only four out of the 24 total cases use the binary version of the response variable, indicating that there **is** something to be gained from transforming the binary response into a multinomial. A third note is that only three out of the 24 cases use the original or un-optimized version of the CART, indicating that optimizing the prior probabilities to increase the RPSS **definitely** leads to improved closed-loop performance.

With dozens of CART models now in our toolset, we can look into the transferability of rules from month-to-month throughout the cooling and swing seasons. Using the closed-loop objective function value as the criteria for selection, the eight best-performing CARTs were chosen, one for each month and one for the aggregate data, and each CART was tested in each of the eight different time periods, resulting in $8 \cdot 8 = 64$ total simulations. Visualizations of the results of each simulation are given in Figure 9.7. The $x - axis$ indicates what data was used to train the CART used in each column, and the $y - axis$ indicates which time-period was used in the simulation. In each of these figures, the results are aggregate sums per-simulation that have been normalized by row, such that the performance of a given CART

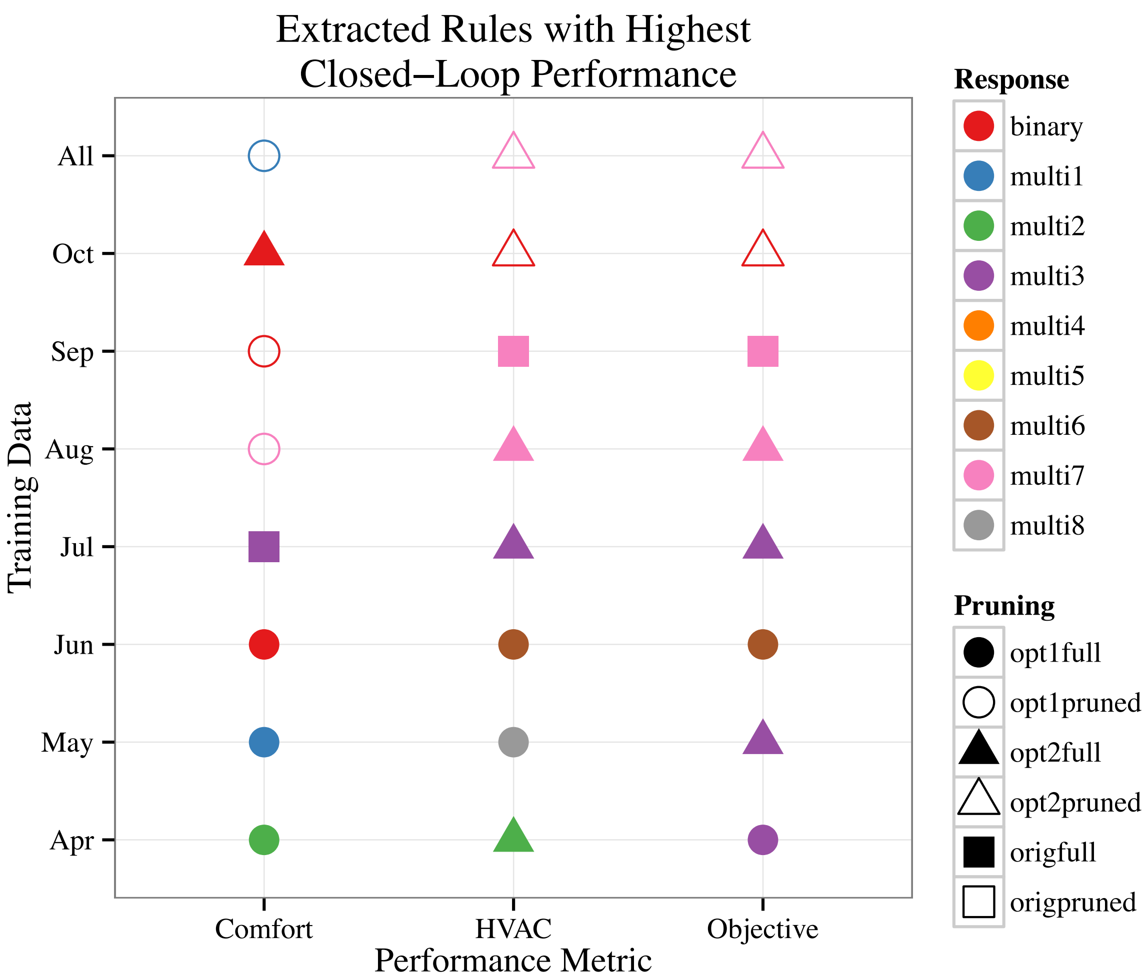

Figure 9.6: Best CART response and growth for monthly rule extraction.

in its own locale (e.g. CART trained on August data, and tested in the August time period) is scaled to a value of 1.0. To keep the color scales of the charts reasonably constrained, any values above 10 (ten times worse than the native CART) are truncated to a value of 10; this is especially helpful with the objective function, which often takes on very large values. Colors in the charts are consistent, such that red is always worse, green is always better, and yellow is acceptable.

In Figure 9.7a, we see what is expected, that the diagonal values are best - indicating that CARTs are seasonally dependent; each rule works best in its own month, and rules trained on other months always perform poorly when tested outside of their training period. One surprise is that several of the one-month rules (Jun, Aug, Oct) outperform the aggregate seven-month rule, when tested on the full seven-month period. Closer inspection reveals that these one-month CARTs tend to leave windows closed more often than the aggregate-CART, which tends to reduce HVAC energy consumption, which we can recall has a stronger influence on the objective function.

One feature we hoped to see in a chart like this and unfortunately do not is good performance of the aggregate rule when tested in each of its constituent monthly datasets. The CART corresponding to September, for example, outperforms the aggregate dataset CART in every month except for May and July (consequently also for the seven-month period). Also notable is a trend that is evident if we look at June and September specifically. Most CARTs trained in other months perform relatively well in June and September, and CARTS trained in June and September tend to perform well in other months, both June and September bridge the transitions between the swing and cooling seasons - and they, like the aggregate dataset, are exposed to a wider range of conditions. On the other end of the spectrum, months in the heart of a given season like April, May, October (swing), and July (cooling), seem to be very specific. With the possible exception of October, these CARTs seem to be particular to their own time period; they do not perform well when tested in other months, and other CARTs do not perform well when tested in April, May, July, or

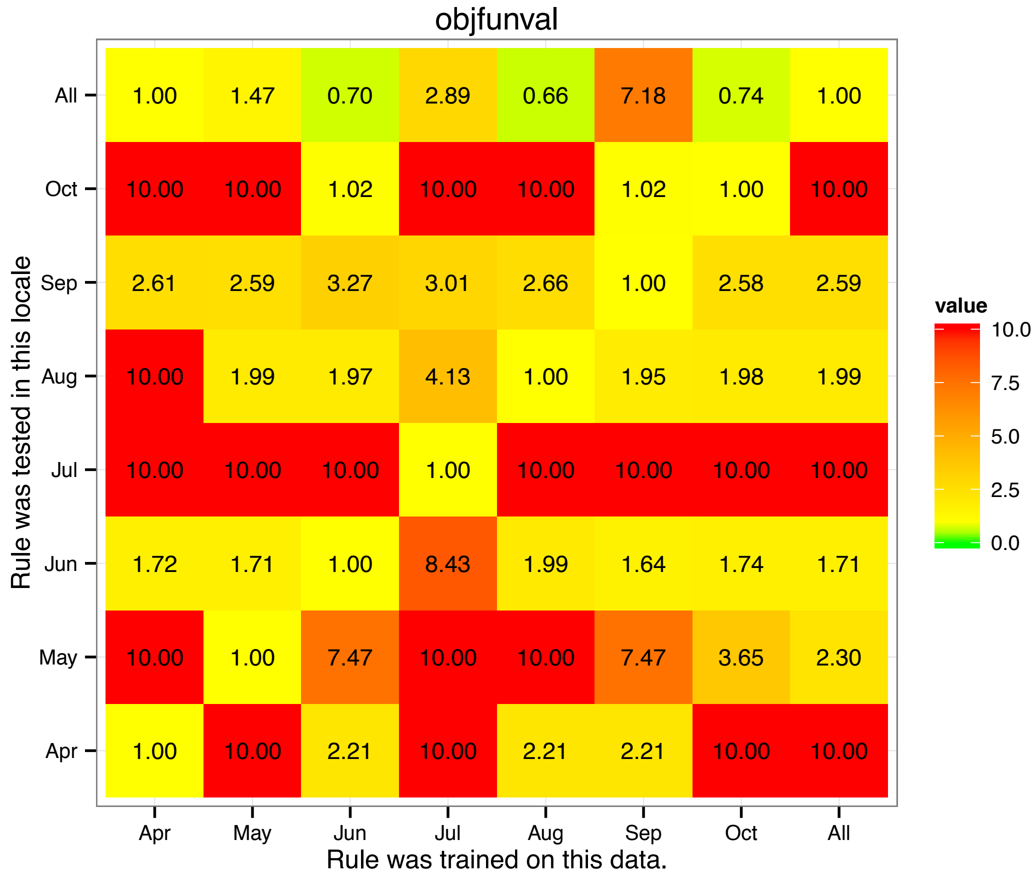

(a) Monthly cross validation: objective function.

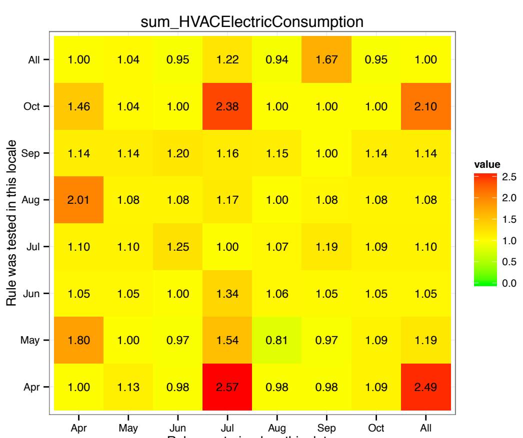

(b) Monthly cross validation: HVAC electric.

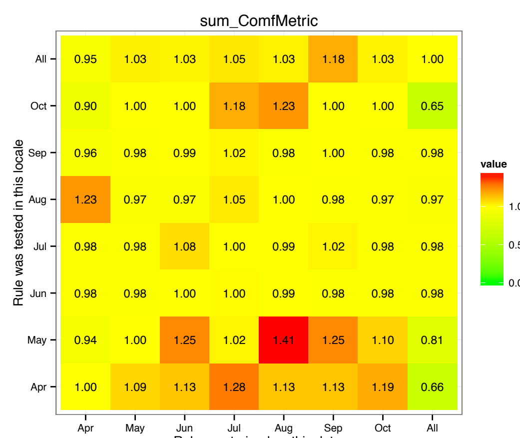

(c) Monthly cross validation: comfort metric.

Figure 9.7: Monthly cross validation: objective function summary.

October.

In summary, it is not clear from looking at the objective function alone that using an aggregate dataset to train a rule on multiple months will ensure good performance in each season. In Figures 9.7b and 9.7c, we can inspect the two components that contribute to the objective function, namely HVAC energy consumption and the comfort metric. The benefit of using an aggregated or larger dataset is more obvious in these charts, since we can see that no rule performs as well in terms of the comfort metric. The CART trained in August does exceedingly well in terms of HVAC energy alone, but struggles in the spring and fall months to maintain comfort.

9.4 Cross-Climate Rule Extraction

Having established that CARTs trained on datasets spanning longer time periods tend to perform better than CARTs trained on individual months, we take a look now at exposing rules to different climates. In the previous section, we trained CARTs on individual months and conducted a cross-comparison of performance; in similar fashion in this section, CARTs are trained based on optimizations for different climates. In total, 42 different locations were selected from across the US to capture at least one location in each ASHRAE climate zone, to capture a number of locations in the vicinity of the actual location of the RSF, in and around Colorado, and in similar climates across the western US.

Table 9.6 lists all of the sites selected. Abbreviations in the Colorado Region column correspond to the Western slope (ws), or western portion of the state, the plains or eastern portion of the state (pl), and the Front Range (fr), locations just East of the rocky mountains. Figures 9.8 and 9.9 show the geographic location of each site in the US and within Colorado, respectively.

After selecting the sites, a deterministic optimization was conducted for each site spanning the time period from April 1 to October 31. Examples of results for two locations, Golden, CO, and Homestead, FL, are given in Figures 9.11 and 9.10, respectively. Note that

Table 9.6: Sites selected for cross-climate rule extraction study.

WMO Number	State	ASHRAE Climate Zone	Site Name	Colorado Region
722026	FL	1a	Homestead AFB	
722050	FL	2a	Orlando Intl AP	
722748	AZ	2b	Casa Grande AWOS	
723110	GA	3a	Athens-Ben Epps AP	
722650	TX	3b	Midland Intl AP	
723910	CA	3c	Point Mugu NAS	
724055	VA	4a	Leesburg Muni AP-Godfrey Field	
724855	NV	4b	Tonopah AP	
725070	RI	5a	Providence-T F Green State AP	
723663	NM	5b	Taos Muni AP	
723677	NM	5b	Las Vegas-Muni AP	
723747	AZ	5b	Show Low Muni AP	
724625	CO	5b	Durango-La Plata County AP	ws
724640	CO	5b	Pueblo Mem AP	fr
724660	CO	5b	Colorado Springs-Peterson Field	fr
724665	CO	5b	Limon Muni AP	pl
724666	CO	5b	Golden-NREL	fr
724698	CO	5b	Akron-Washington County AP	pl
724699	CO	5b	Boulder-Broomfield-Jefferson County AP	fr
724756	UT	5b	Bryce Canyon AP	
724760	CO	5b	Grand Junction-Walker Field	ws
724765	CO	5b	Montrose County AP	ws
724767	CO	5b	Cortez-Montezuma County AP	ws
724768	CO	5b	Greeley-Weld County AWOS	fr
724769	CO	5b	Fort Collins AWOS	fr
724860	NV	5b	Ely-Yelland Field	
725717	CO	5b	Rifle-Garfield County Rgnl AP	ws
725724	UT	5b	Provo Muni AWOS	
725805	NV	5b	Lovelock-Derby Field	
725866	ID	5b	Twin Falls-Magic Valley Rgnl AP-Joslin Field	
725955	CA	5b	Montague-Siskiyou County AP	
726835	OR	5b	Redmond-Roberts Field	
726886	OR	5b	Baker Muni AP	
726988	WA	5b	The Dalles Muni AP	
727810	WA	5b	Yakima Air Terminal-McAllister Field	
727830	ID	5b	Lewiston-Nez Perce County AP	
727834	ID	5b	Coeur dAlene AWOS	
725945	CA	5c	Arcata AP	
727885	WA	5c	Port Angeles-William R Fairchild Intl AP	
726380	MI	6a	Houghton-Lake Roscommon County AP	
726776	MT	6b	Lewistown Muni AP	
727477	MN	7	Roseau Muni AWOS	

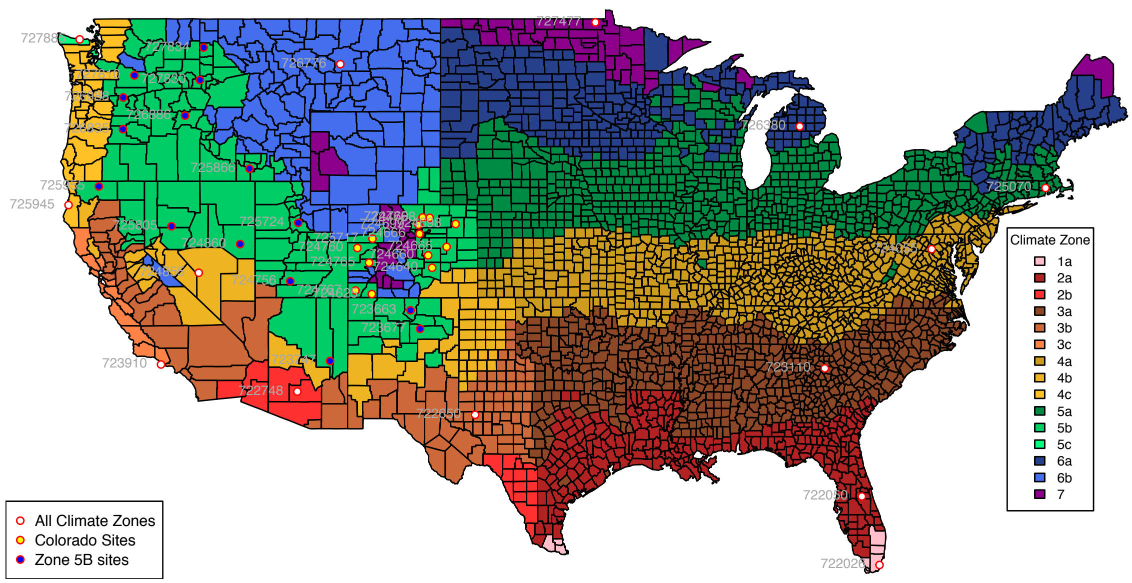

Figure 9.8: All US-WMO sites selected for cross-climate rule extraction study.

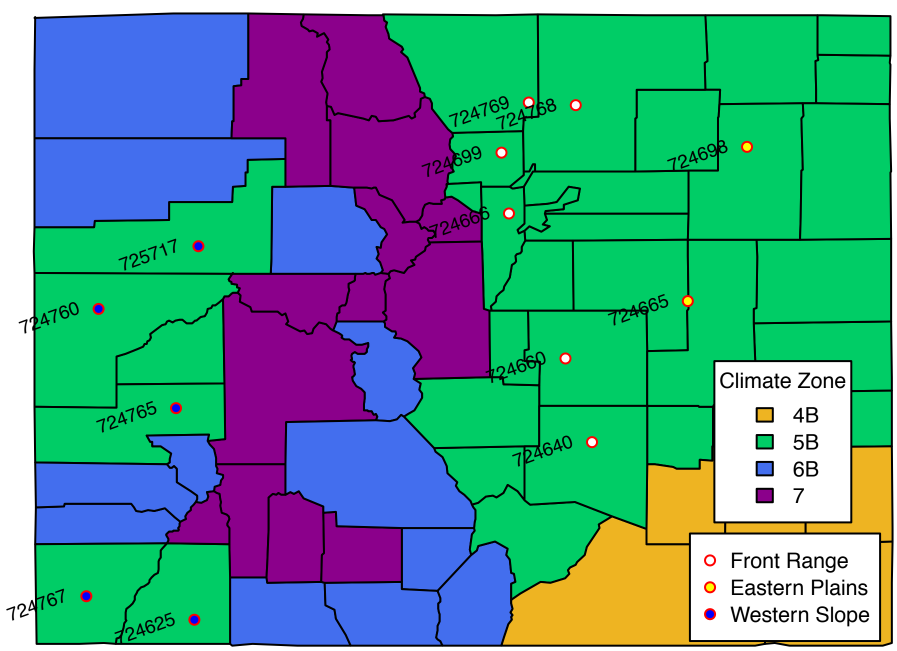

Figure 9.9: Colorado WMO sites selected for cross-climate rule extraction study.

energy savings are nearly impossible in the Florida location (climate zone 1a: Hot and Humid), however comfort is improved by occasional window openings. CARTs generated from the results of the Homestead, FL optimization depend heavily on indoor and outdoor humidity levels, while those generated from Golden, CO results use indoor CO_2 concentration and indoor temperature to control natural ventilation.

In all of the 42 optimizations conducted, total HVAC savings ranged from 1-19% over the seven-month period, and primarily fell in the 5-10% range, indicating that the RSF building is (as designed) extremely efficient, and performance improvements are difficult in any climate. Improvements in the comfort metric were negligible, with 5% improvement in a handful of cases.

Just as Figure 9.6 showed the best rules for each monthly CART, Figure 9.12 shows the best CART for each climate. In this figure, we begin to see definite trends; multinomial models dominate the group, as do full (un-pruned) models. Multinomial-7, -3, and -1 models occur frequently, but the Multinomial-2 format seems to be the best overall, counting all three metrics (HVAC energy alone, comfort alone, and the objective function) and all 138 total cases, the Multinomial-2 case scores best 44 times. Recall from Table 9.3 that Multinomial-2 corresponds to transforming a binary signal to a 3-level multinomial signal, where the change in state is associated with the hour preceding a change. This makes sense given that we want the CART to leverage some predictive power, if possible, just as the optimizer does in MPC. The six additional cases in the top of Figure 9.12 correspond to CARTs trained with six different aggregate datasets. Since optimizing the natural ventilation controls of the RSF was the primary goal, the six aggregate datasets, and the selection of the 42 climate sites followed the logic that rules would be more robust when trained with more data. To get a larger dataset, the idea is to optimize the RSF building in multiple climates; the first aggregate dataset, Colorado: Front Range, includes six sites, all within a few hundred miles of the RSFs home location in Golden, CO. All of these sites are located in the transitional region between the Rocky Mountains and the flat eastern plains of Colorado, which is a

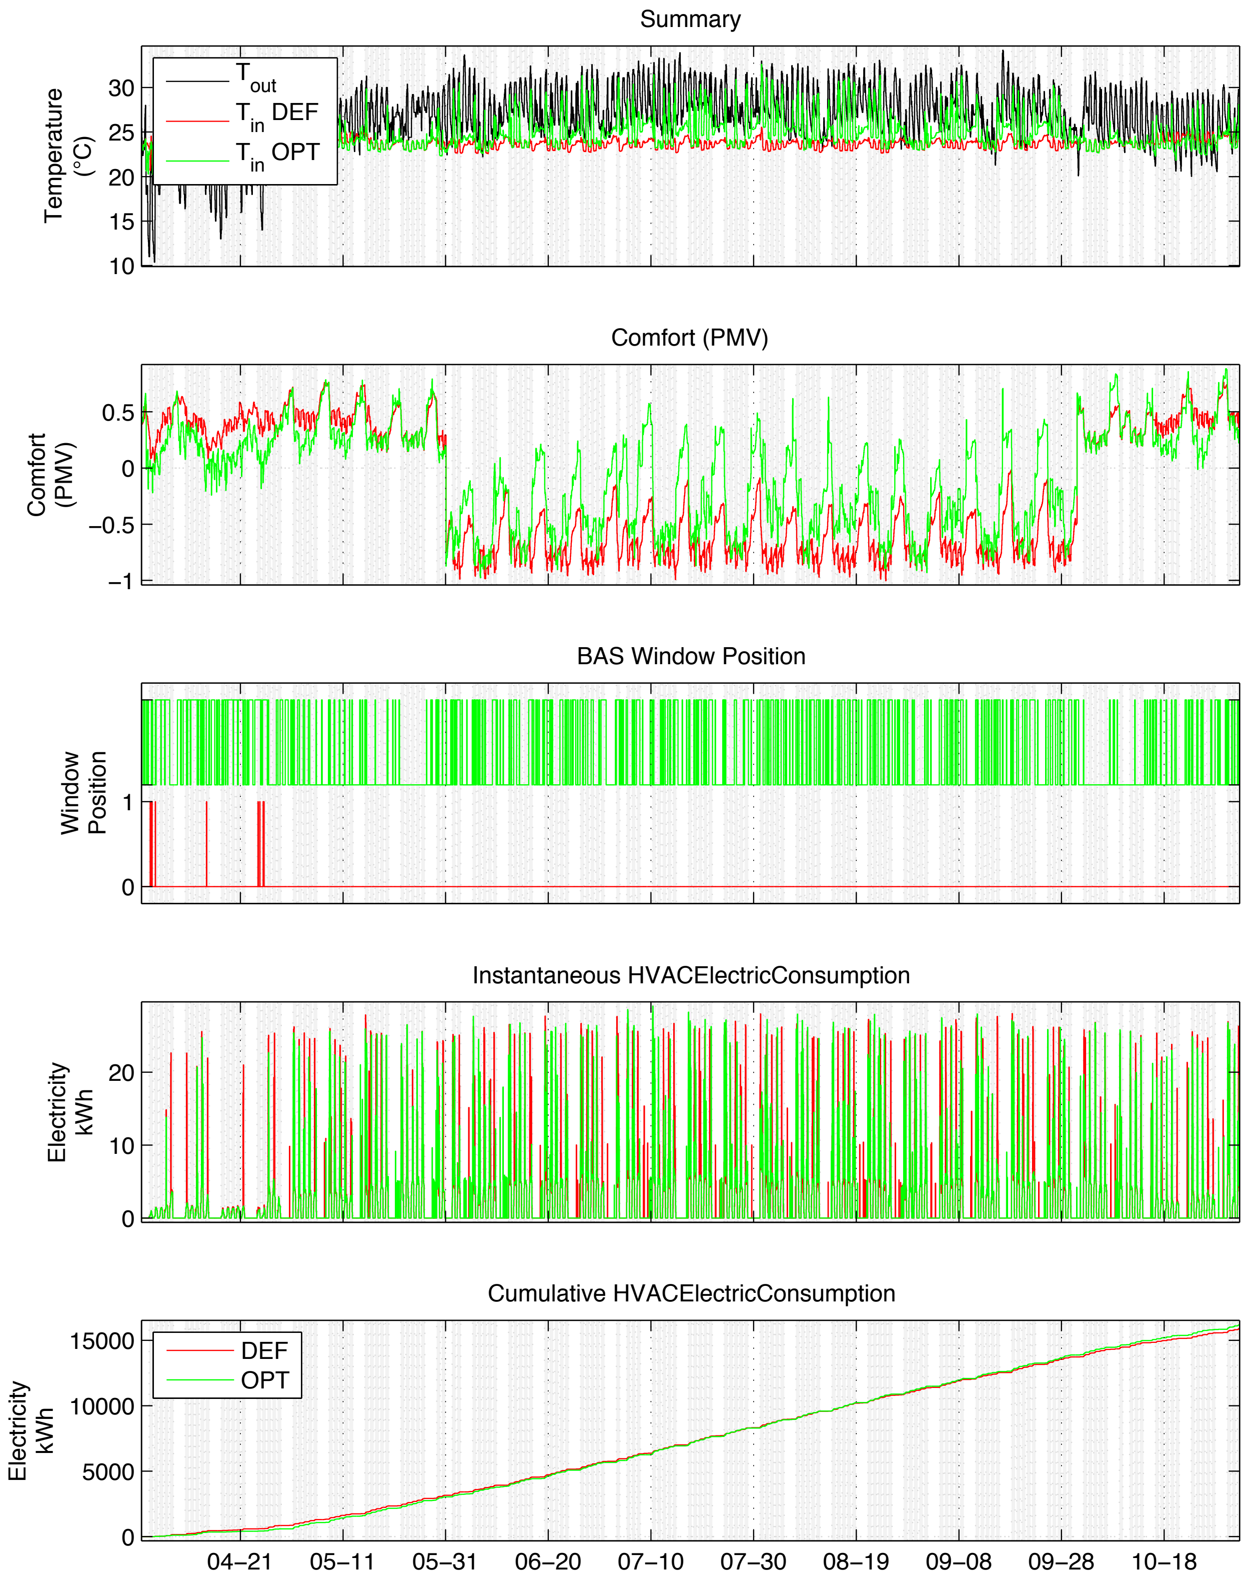

Figure 9.10: Florida climate zone 1a optimization results summary.

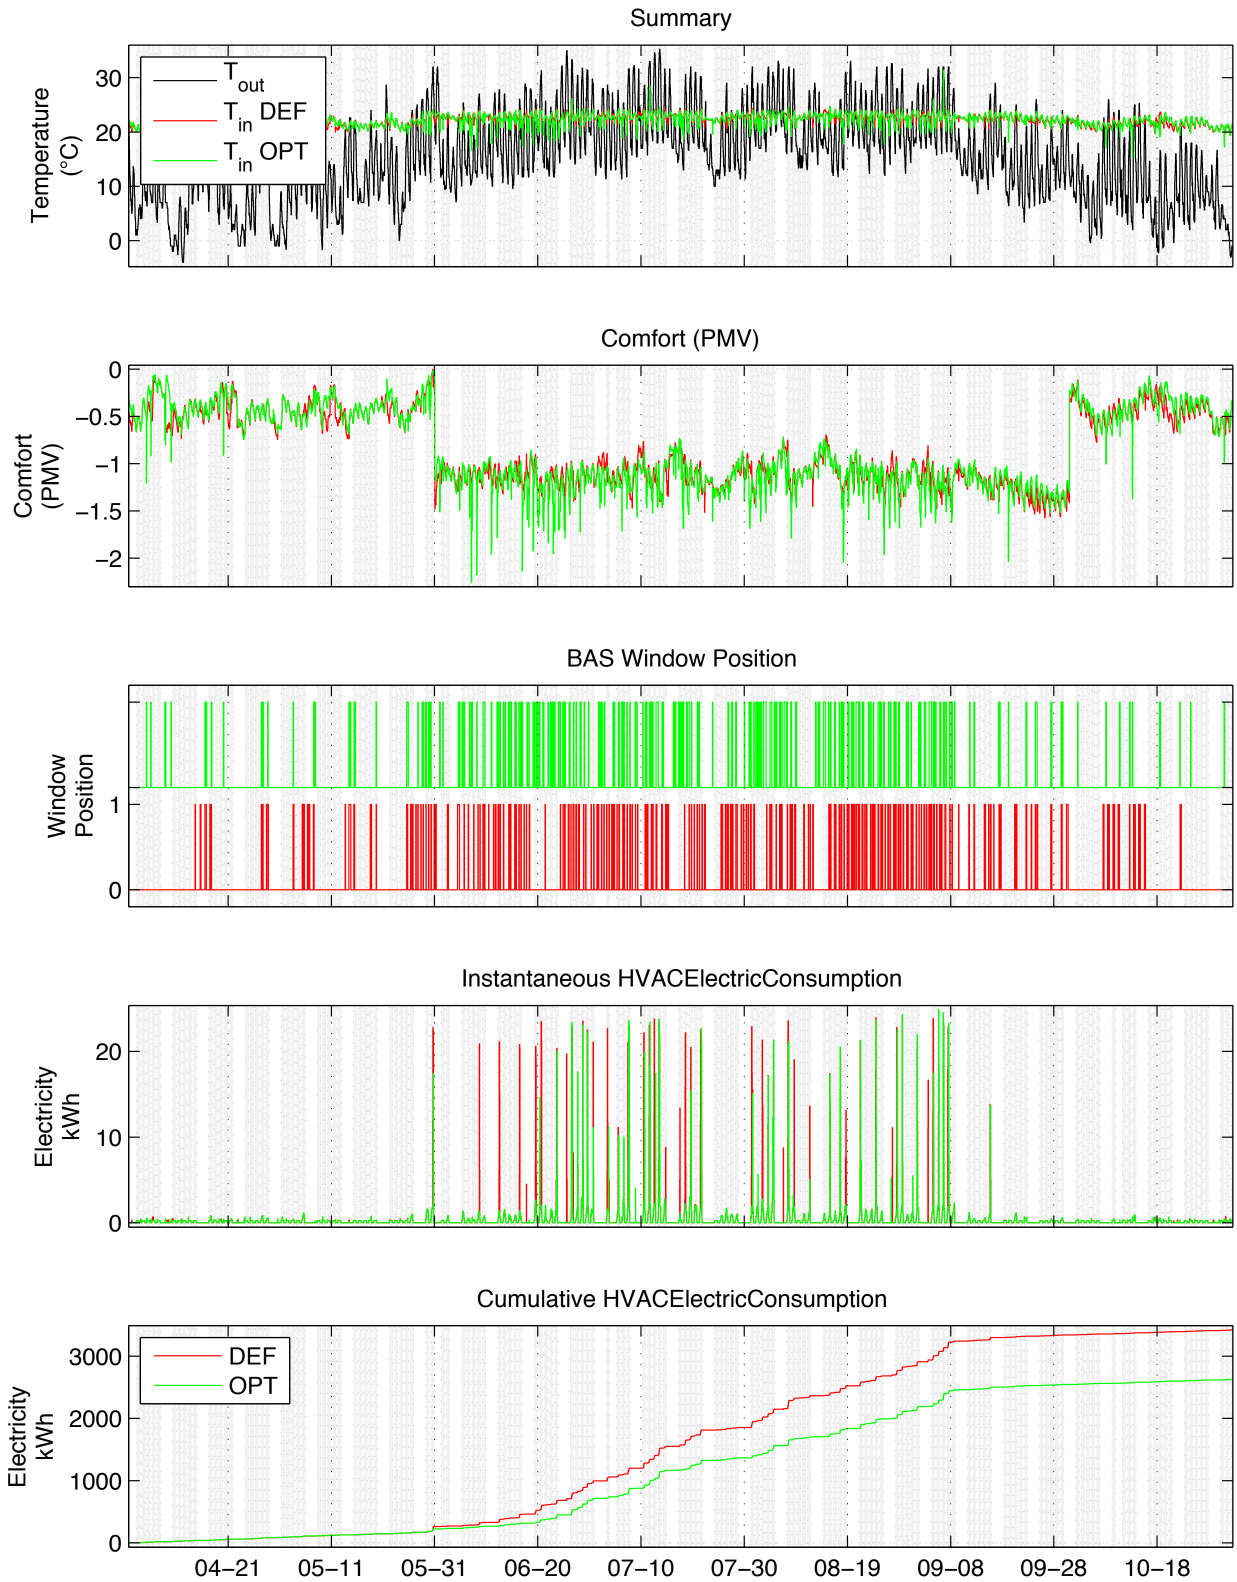

Figure 9.11: Golden climate zone 5b optimization results summary.

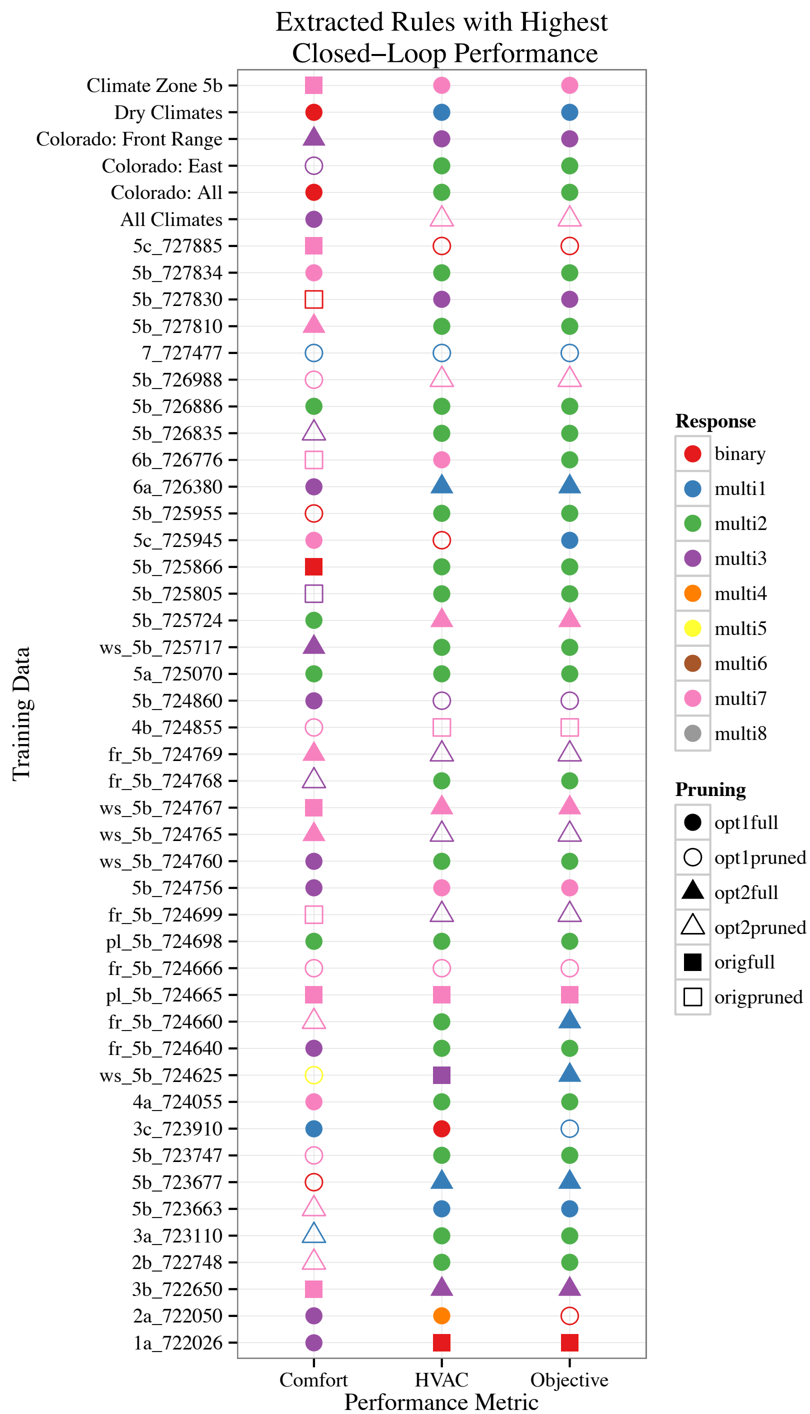

Figure 9.12: Best CART response and growth for cross climate rule extraction.

region with very specific weather patterns - significantly different from either the nearby mountains or the nearby plains.

The second aggregate dataset, Colorado: East, consists of the same six sites in the Front Range group, as well as two more that are in the nearby easter plains of colorado. The third set, Colorado: All, includes five sites located in the western part of colorado that certainly experiences different weather patterns - but still falls in the ASHRAE 5B climate zone. The fourth set expands on the third by including 15 additional climate zone 5B sites from across the western US. The fifth set adds in four locations that are in dry climate zones (2B, 3B, 4B, 6B), and the final set includes nine more sites to include at least one site from each of the ASHRAE climate zones represented in the continental US. By exposing the RSF building model to a variety of climates in simulation, with an emphasis on nearby or similar locations, we hope to train a CART that is robust to **any** environmental conditions that might come along. By choosing the six different sets of aggregate data, we also hope to tease out what level of aggregation - or how much variety - is necessary to train a highly robust CART.

We begin breaking down the results in Figure 9.7 with a look at the the extremes. Locations 722026 and 727885 correspond to climate zones 1a and 5c respectively; one is on the southern tip of Florida and the other is on the Olympic peninsula in Washington State. Understandably, these locations represent the edge-cases and rules for these locations are very specific to the local climate. Not only do these CARTs not perform well anywhere else, but almost no other rules are capable of performing well in these locations.

Next we can take a look at our primary location of interest, site 724666 (Golden, CO), which seems to be unique in that it is a relatively difficult climate to handle for CARTs trained elsewhere, and the CART trained on Golden, CO weather data only works well in its own location. Interestingly, there is not one other location with both of these characteristics.

Next, we observe three sites which are difficult locations for CARTs trained elsewhere to work in, but tend to generate CARTs that perform well everywhere else. Sites 724756, 724767, and 725724 correspond to Bryce Canyon, UT, Cortez, CO, and Provo, UT, respectively. All three locations are in climate zone 5B, and all three CART models use the Multinomial-7 response format, with fully grown trees. The Bryce Canyon CART uses an original CART, while the other two use the second growth optimization.

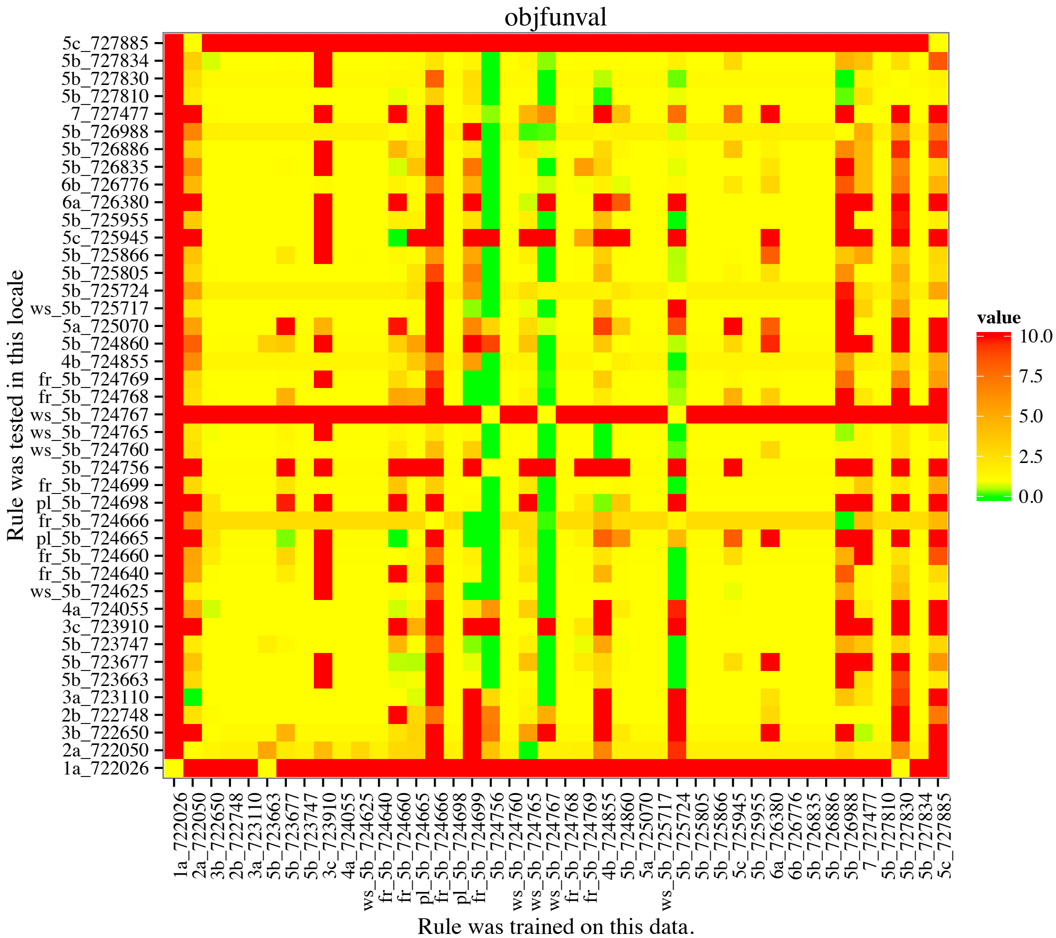

(a) Climate cross validation: objective function value.

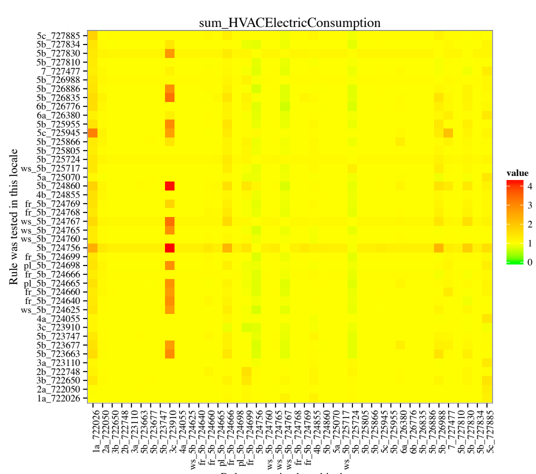

(b) Climate cross validation: HVAC electric.

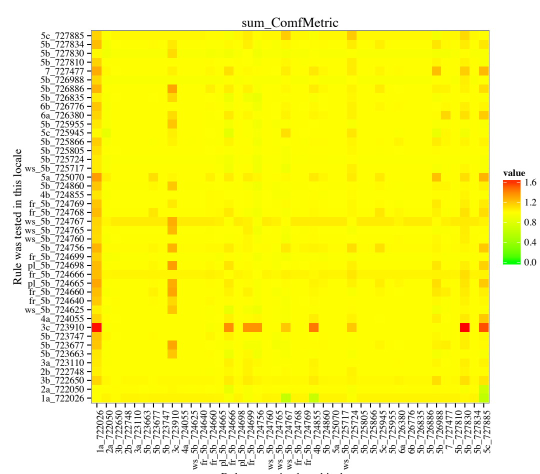

(c) Climate cross validation: comfort metric.

Figure 9.13: Climate cross validation: objective function summary.

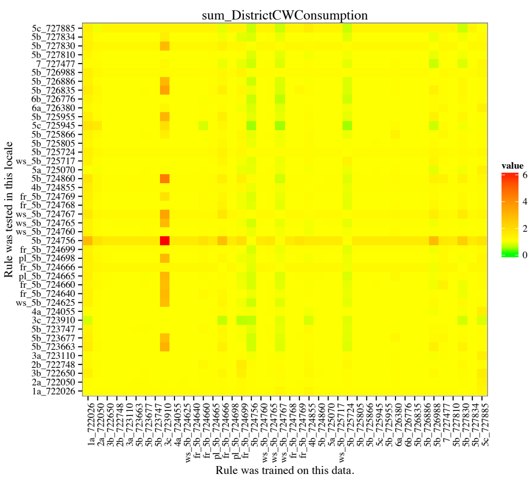

(a) Climate cross validation: cooling.

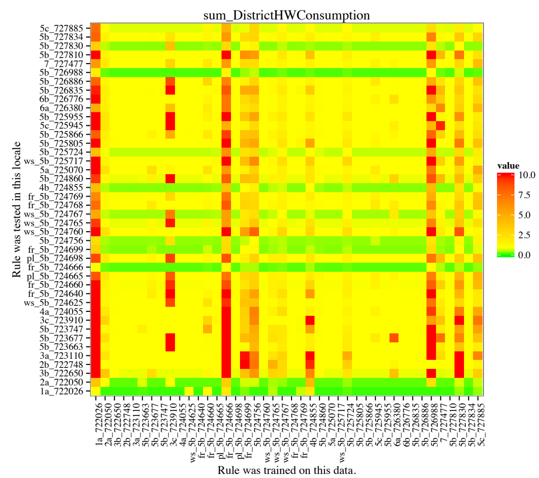

(b) Climate cross validation: heating.

Figure 9.14: Climate cross validation: heating and cooling.

Chapter 10

RSF Controls Implementation

10.1 Introduction

The RSF building is the largest net-zero energy buildings in the world, it operates at a site energy use intensity of 25 kBtu/ft^2, and a comfort study conducted as part of this project ranked the RSF in the top 92nd percentile for comfort satisfaction among occupants in 2012, and in the 95th percentile in 2013. Of the energy consumed in the RSF, 3% is cooling-related, and 6% goes to powering fans and pumps. This chapter addresses the process of attempting to improve comfort and/or reduce energy consumption from fans, pumps, and cooling in the RSF building. Achieving meaningful improvements in any of these areas is rather difficult, given the existing near-optimal performance of the building.

10.2 Summary

Our original goal was to use stochastic model predictive control to generate a dataset of optimal control decisions for the automatic windows in the RSF, and to use this dataset and rule extraction methods to generate simplified near-optimal control rules that we could implement in the building - proving the methods can work outside of simulation.

An initial meeting with RSF personnel in Spring, 2012 led to an informal agreement to use the RSF building as a testbed later that summer and/or in the warmer months of 2013. The RSF staff generously provided a complete energy model of the RSF building, however the model was deemed unusable since it took 24 hours to run a single annual simulation, and it did not incorporate radiant systems, underfloor air distribution, or natural ventilation - the three primary systems involved in the cooling and ventilation of the building. The model did incorporate work-arounds which enabled the model to provide estimates of energy consumption throughout the design process.

Work commenced with constructing the single zone RSF model, which was calibrated to initial annual energy usage statistics. Next, occupant behavior models were coded into the building simulation model, specifically the Haldi model for occupant use of windows [38].

Results of initial parametric studies with the single-zone model provided preliminary insight into savings opportunities in the RSF. Savings from adjusting only the setpoints in the existing control logic were small, and work continued with SMPC investigations in the fall of 2012.

In mid summer 2013, the decision was made to abandon the top-down approach in favor of the simpler bottom-up approach so that a controls change could be made before the cold weather set in later in the fall, ending the natural ventilation season.

Additionally, making involved changes to control logic statements or programming new algorithms was not possible due to budgetary and manpower constraints within NREL's facility management department, so a decision was made to focus on optimizing a few setpoints in the natural ventilation control logic, similar to the parametric study results presented in Chapter 7.

10.3 Controls Changes

Using results of past parametric studies, optimizations, and results of new parametric studies conducted in the summer of 2013, changes in setpoints were recommended that govern automatic window openings and closings, as well as the setpoints that govern variable air volume (VAV) terminal positions when windows are open. A summary of the various RSF studies, and the final control changes made are given in the following:

Parametric Study 1 This parametric study is described in detail in Chapter 7, and took place in June of 2012, and focused on setpoints that govern night-ventilation. We investigated the impact of changing the minimum outdoor air temperature setpoint that governed automatic window openings, as well as the time of day when night ventilation was allowed. At the time, the minimum setpoint was 65°F, and the night ventilation was enabled between 6pm and 6am. Our study showed that 65°F was the best choice (between 61 and 69°F), and that night ventilation had the largest benefit in the morning - though enabling it at 6pm did not result in significant comfort or energy penalties.

Parametric Study 2 In the summer of 2012, independent of our work, the site operations and engineering staff at the RSF resolved to enable natural ventilation at all times, not just at night. Given the new controls, we conducted a new parametric study, again looking at minimum outdoor air temperature, but this time also considering adding window-HVAC interlock. The interlock system would simply close VAV terminal boxes whenever automatic windows were open, rather than keeping VAV terminal boxes open at all times, regardless of natural ventilation - as was the case at the time. The outcome of this study were multiple; the major insight was that the addition of an HVAC-interlock system would generate the most energy savings, independent of all other natural ventilation controls. Secondary results were that the majority of potential energy savings in the RSF were available during the swing seasons and that a seasonal or monthly reset of setpoints (minimum outdoor air temperatures) would lead to better performance than a single annual setpoint.

Stochastic MPC and Offline Optimizations Results of our optimizations with the RSF have shown that there is minimal room for improvement in energy performance at the RSF from natural ventilation due to the already outstanding performance of the building. The RSF was designed to use a small amount of cooling energy, and so reducing cooling energy consumption is challenging. Our optimizations have nevertheless led us to a few opportunities for savings. The first opportunity is one common to most mixed mode buildings, the opportunity for pre-cooling; charging the building's thermal mass with cool outside air before a warm day when the building will require cooling. As in the parametric studies, our optimizations have shown that the swing months have the largest potential for generating energy savings. The final lesson learned from our optimizations is that in the spring, when the building is switching from primarily heating to primarily cooling (or vice-versa, in the fall), it is important to avoid switching back and forth multiple times, since the building has to heat up or cool down a massive concrete ceiling. Natural ventilation can help to avoid repeated switching between heating and cooling by providing natural cooling at times during these swing months - instead of forcing the radiant systems to fully cool the concrete ceiling.

Parametric Study 3 In the summer of 2013, with the results of the last three studies in mind and a new directive from the engineering staff at the RSF, we worked through a final parametric study. This final study was the result of several meetings and brainstorming sessions which leveraged the results of the previous studies and optimizations to design new simplified window control algorithms. The goal was to try five or six different pragmatic control strategies, to look in detail at where they performed better or worse than the existing controls, and then to see how much benefit could be gained by changing only setpoints in the existing controls. The different strategies employed in this study were:

Base Case (Existing Controls) Windows are opened whenever wind is minimal, outdoor relative humidity is below 50%, and outdoor temperature is between 68 and 74°F. Windows are closed whenever wind is gusty or strong, outdoor relative humidity is above 52%, or outdoor temperatures fall outside of the

68-74°F range.

Case 2: Indoor Temp vs. Heating Setpoint Windows are opened according to the rules in the Base Case. Windows are closed whenever **indoor** temperature approaches the heating setpoint. (intended to allow for more cooling than the base case)

Case 3: Indoor vs. Outdoor Temp Windows are opened according to the rules in the Base Case. Windows are closed whenever **indoor** temperature approaches the outdoor temperature (this is intended to allow aggressive cooling).

Case 4: Higher Humidity Windows are opened and closed according to the rules in the Base Case, with the exception that the humidity setpoint is not 50%, instead multiple options for relative humidity were selected (60, 70, 80%).

Case 5: Relative Enthalpy Windows are opened whenever outdoor enthalpy is lower than indoor enthalpy, and temperatures are cooler than 74°F. Windows are closed whenever outdoor enthalpy approaches indoor enthalpy, or outdoor temperature drops below a minimum (60°F).

Case 6: Absolute Enthalpy Windows are opened whenever outdoor enthalpy is lower than a fixed upper limit, equal to the enthalpy at 74°F and 50% RH (the high limit in the base case). Windows are closed whenever outdoor enthalpy approaches the upper limit, or outdoor temperature drops below a minimum (60°F).

The outcome of the final parametric study (beyond reinforcing the conclusions made in prior studies) was a conclusion that little could be gained from changing setpoints in the existing controls, and that truly a comparison with indoor conditions was necessary to make a meaningful improvement.

Final Controls Recommendations Leveraging the results of all of the controls investigations listed above, our final controls recommendations (sent on September 6, 2013) to the engineering staff were as follows.

(1) The controls change that will have the most significant impact is to enable VAV terminal boxes to close fully whenever automatic windows are open.

(2) The second most meaningful change that can be made to the existing window controls is to reduce the minimum outdoor air temperature from 68 to 64°F.

(3) If changing the window control logic is possible, we recommend the following logic: open windows whenever **outdoor** conditions are satisfactory (as in the existing controls), but close windows whenever **indoor** conditions are approaching unsatisfactory levels.

Final Controls Changes Implemented In Short:

(1) Minimum outdoor air temperature setpoint for closing windows reduced from 68 to 65°F.

(2) Three VAV terminal boxes have lower minimum airflow setpoints.

In detail: In response to our recommendations, the staff of the RSF agreed to reduce the minimum outdoor air temperature setpoint for **closing** windows to 65°F, since they had used that value previously without major consequences. The setpoint for **opening** windows remains at 68°F. This is a slight deviation from our recommendation, which was to both open **and** close windows around the 64°F setpoint, but opening windows when temperatures were that cool seemed too aggressive for the RSF staff.

The RSF staff was however unable to fully implement the other controls changes for different reasons. Fully closing VAV terminal boxes when windows are open was ruled out, since it would require adding logic to the VAV controls programs (i.e.

checking for a natural ventilation status, and adjusting VAV terminal boxes accordingly), however reducing the minimum damper setpoint was a feasible change. As of November 2013, three VAV terminal boxes have had minimum airflow setpoints reduced from 130, 140, and 135 to 80, 90, and 80, respectively. Results presented below show that reducing the airflow through these VAV terminal boxes has not resulted in larger CO_2 concentrations than were present with the larger airflow setpoints.

In aggregate, the airflow reductions for the three VAV boxes represent an 8% reduction in the total airflow through the main AHU for that wing of the building; if all VAVs in the wing were changed proportionally, the main AHU would see a 35% reduction in airflow.

Fully changing the window control logic to respond to indoor conditions was ruled out because it would require unique code for each zone of the RSF. Currently all of the windows are controlled at once for 14 different zones, and making 14 individual programs (each computing local indoor average/minimum/maximum temperatures and relative humidities) was not possible; however the site operations team has taken the suggestion under consideration and may implement it in the future.

10.4 Recommendations

What can a practitioner learn from this experience? In the RSF the granularity of window control is limited to fairly large portions of the building. For the 14 zones mentioned above, there could be as many as eight natural ventilation zones based on the window actuating hardware that is in place. The bottom floors are grouped together into one control zone for each wing, and the top floor exists as its own zone for each wing (there are four wings total in the original RSF building; a two-wing addition is not included in this analysis). So at most, the controls can be refined to meet the needs of eight different natural ventilation zones individually, but each of these zones can see 2-3 °F variations in temperature throughout the space. Parts of the each space include small meeting rooms, kitchen areas, and copy areas, and it is challenging to pick one representative temperature sensor to guide controls for natural ventilation for these different spaces that are all grouped together. This is a large part of the reason that automatic windows are only controlled according to outdoor conditions - because the indoor conditions can vary widely within an individual natural ventilation zone.

That said, it would likely benefit the RSF to break up the natural ventilation controls into the eight possible groups, since the lower floors experience different envelope loads than the top floor, and the the east wings experience different loads than the west wings over the course of each day. The ground floor in an east wing is controlled exactly the same as the top floor in a west wing, but on a given temperate afternoon it might be too cool for natural ventilation in the east-ground floor space, but perfect for natural ventilation in the west-top floor space.

For natural ventilation designers, we suggest the following conventions when designing a natural ventilation system with automatic windows. Specify that natural ventilation systems are zoned according to building space types, and specify that controls for automatic windows be as flexible and localized as possible. If there will be individual offices, then each office should include an individual controller. In general, the bottom floor of a building should be controlled individually because it can represent a security threat (open windows at night could admit burglars), and the top floor should be controlled individually due to the additional envelope heat transfer through the roof. In addition to keeping window actuation hardware and controls as localized as possible, it is important to specify what outdoor **and** indoor sensors will be required to guide window control, and where they should be located (close to the window they are controlling).

Finally, in order for a natural ventilation system to save energy in a building, it needs to be controlled in concert with the mechanical ventilation and/or mechanical cooling systems. If saving mechanical ventilation energy is the goal, then the mechanical ventilation system should ramp down or turn off when natural ventilation is in effect. If saving cooling energy is the goal, then at a minimum, automatic windows need to be opened when outdoor conditions are cooler than indoor conditions.

10.5 Final Comments

In the last 18 months working with the RSF, we have learned just how difficult it is to implement a meaningful change to the controls of a building, largely because of the risk-averse nature of facilities management staff, and also because of the difficulty in accurately predicting building performance from a model. Model creation and calibration are not trivial tasks, and require considerable time and effort. Vetting the results of the model, and convincing ones self and others that what works in the simulation will work in reality, are also difficult tasks.

In the end, we have fallen short of our lofty goal of implementing control rules that were learned from offline optimization results - but given the difficulty we faced in getting even the simplest setpoint changes made, that goal was clearly unreachable. We are confident that the controls changes that have been made have led to some energy savings in the RSF: additional free cooling from leaving windows open at cooler temperatures, and fan savings from reducing airflow through VAV terminal boxes.

Figure 10.1 below shows the status of the three VAV terminal boxes before and after the controls change. Note the immediate drop in airflow after the controls change. Generally speaking, these VAV terminal boxes operate at their minimum values at all times, unless a large number of occupants are breathing near the CO_2 sensor, in which case the VAV terminal box opens to allow more fresh air in to drive CO_2 values down. It remains to be seen whether the new minimum value is still higher than necessary, or if it has been reduced to a point that the VAV terminal boxes actually open and close in response to higher and lower occupancy levels.

Figure 10.2 shows the behavior of the automatic windows after the controls changes

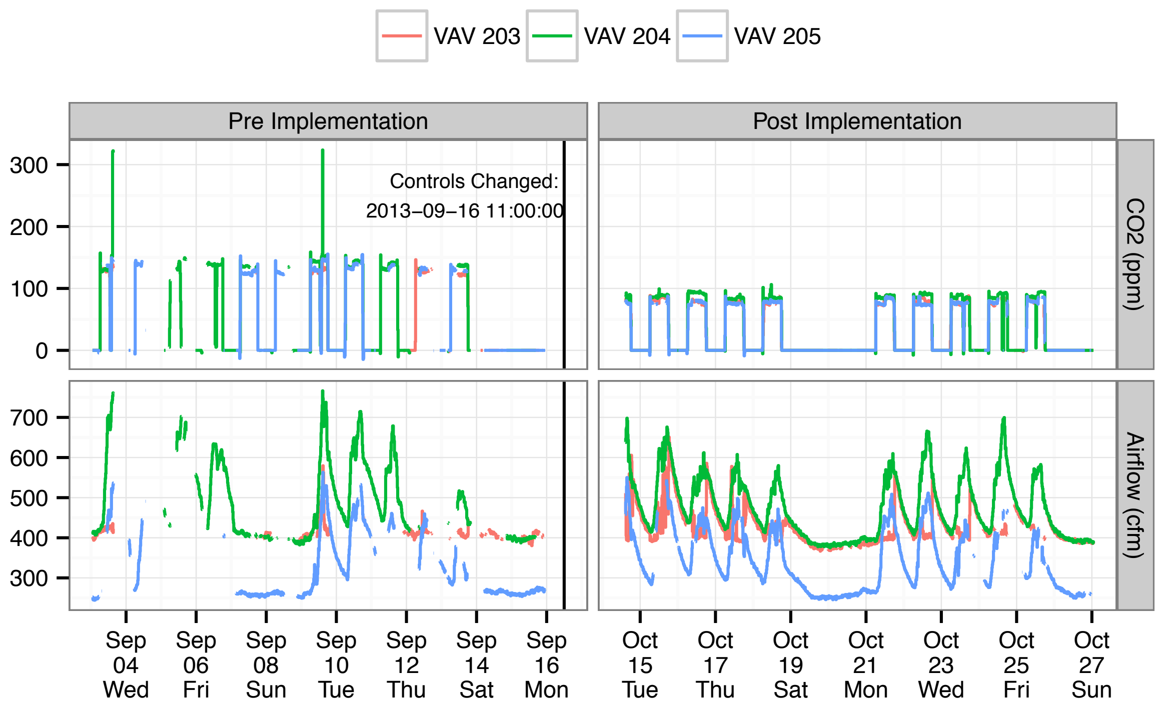

Figure 10.1: Airflow and CO_2 concentration for three VAV terminal boxes in the RSF. Note that the data collection system was down for roughly one month from Sep 17 - Oct 14, and blank data has been omitted.

were implemented. In this case, the controls were changed for RSF1 (wings B and C), but not for RSF2 (wing A), so while the new control sequence for B and C wing allowed windows to stay open until the outdoor air temperature dropped to 65°F, the old setpoint forced windows to close when the outdoor air temperature dropped to 68°F.

Figure 10.3 shows three operational parameters for the air handling unit (AHU) which serves the three VAV terminal boxes that now have lower airflow rates. The combined flow rate through the three VAV terminal boxes accounted for 22% of the total airflow through the main AHU prior to the reduced flow rate setpoints. The total reduction in flow is equivalent to 8% of the total flow through the main AHU, and we expect to see this reduction in the average flow rate during occupied hours, but there is no clear reduction evident in the figure.

Another look at the pre- and post-implementation data from the air handling unit in Figure 10.4 shows that there is not a significant change at the AHU, in spite of the 8% decrease in air flow rate.

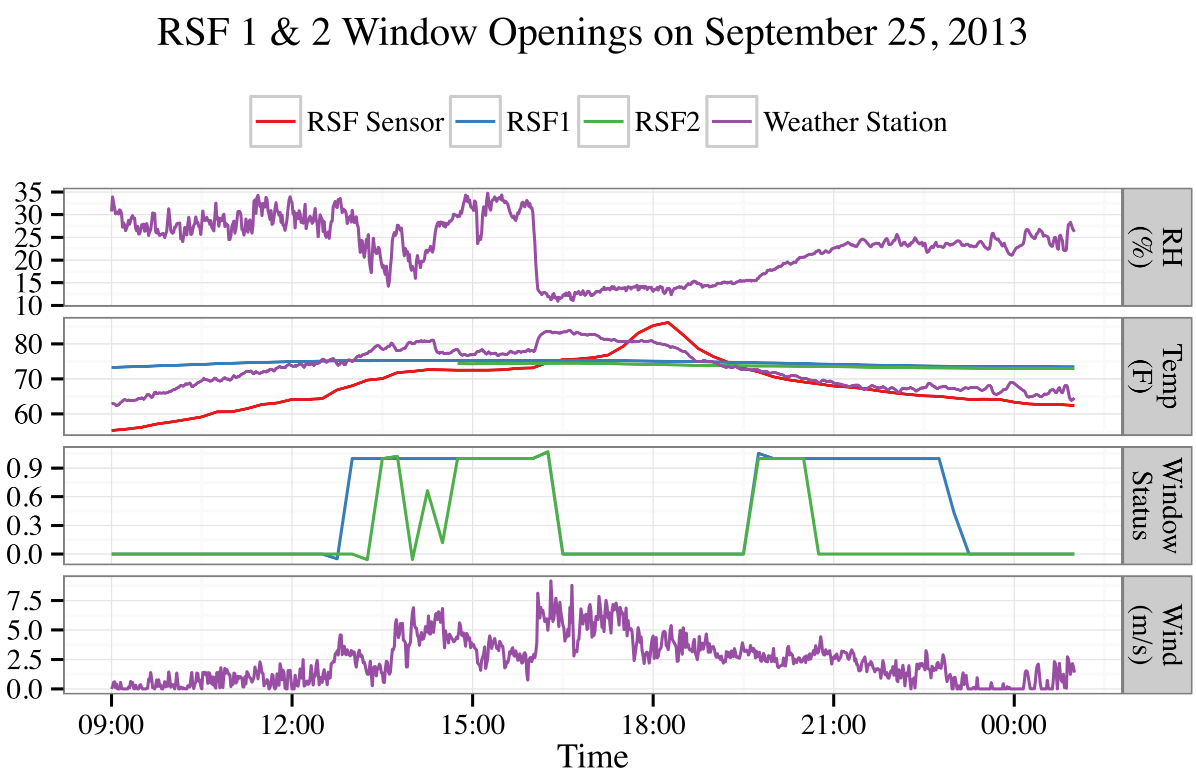

Figure 10.2: RSF window operation after minimum outdoor air temperature threshold for closing windows was reduced to 65°F. Note that the RSF 2 wing of the building retained a setpoint of 68°F, and closed roughly two hours earlier.

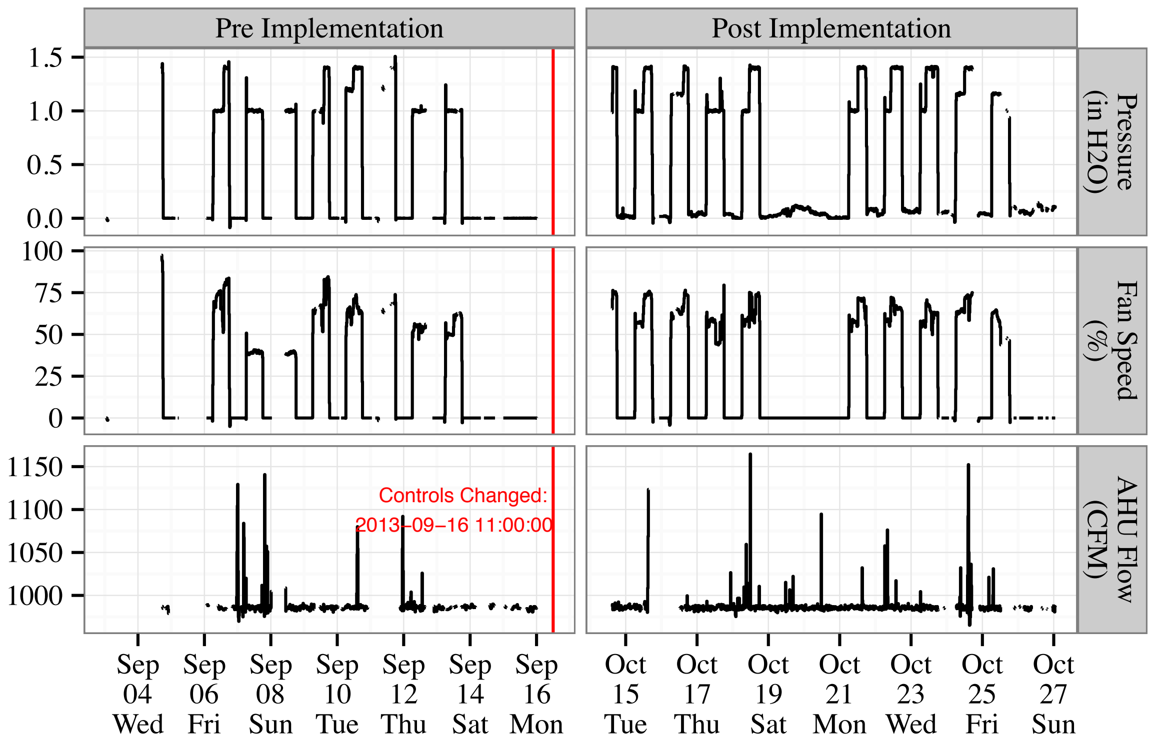

Figure 10.3: RSF air handling unit N1: primary operational parameters before and after minimum flow rates were reduced in VAV terminal boxes 203, 204, and 205.

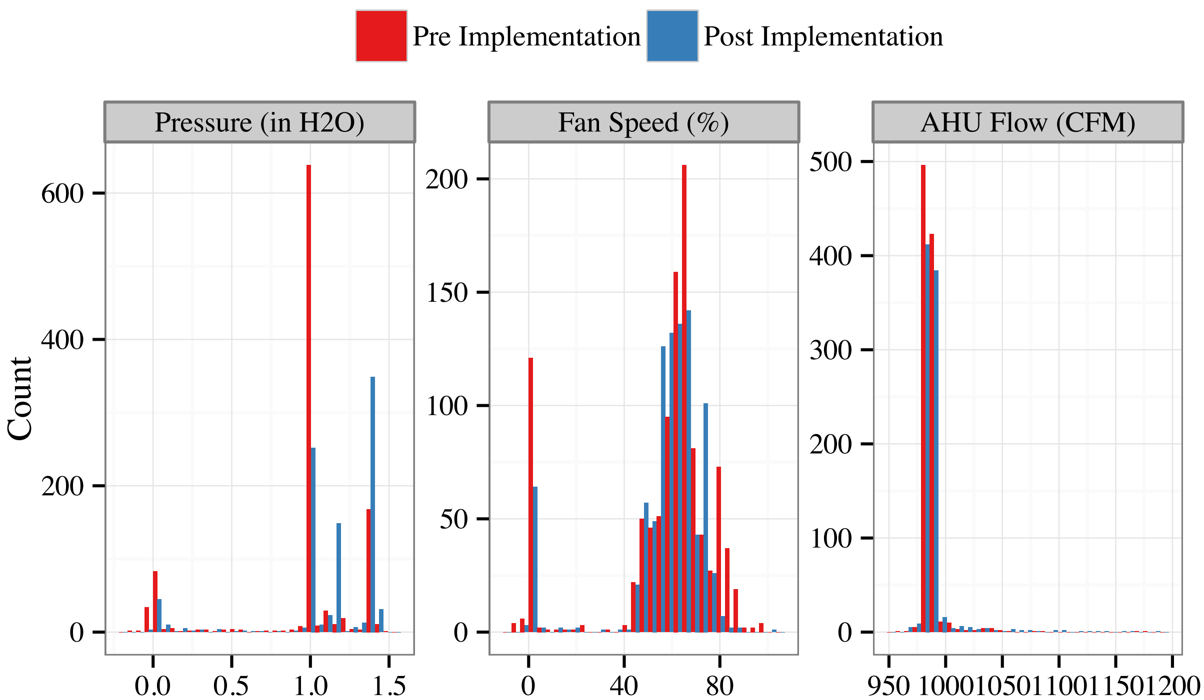

Figure 10.4: Data for occupied (6:00am-6:00pm) hours only.

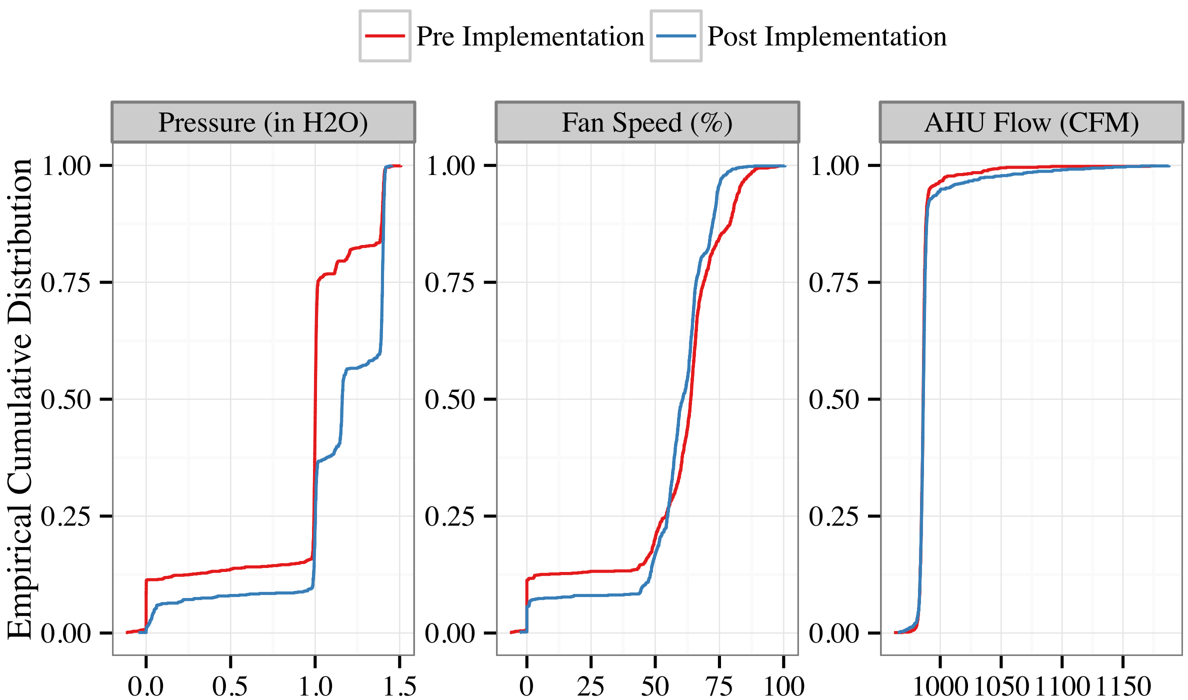

Figure 10.5: Data for occupied (6:00am-6:00pm) hours only.

Chapter 11

Summary, Conclusions, and Future Work

The work presented here was motivated by the need to reduce energy consumption in commercial buildings. To address this need, we focused on mixed mode buildings as a promising means for reducing building energy demand. Recognizing that MM buildings are challenging to control and that occupant actions can significantly impact MM building performance, we built on existing work to investigate opportunities for improving MM control strategies while accounting for the actions of building occupants.

11.1 Summary

From a review of relevant literature, we learned that occupant behavior modeling is a field of growing interest, and that when models of occupant behavior are implemented in building simulations, the impacts on building performance vary significantly, providing valuable information on the robustness of a building to the actions of occupants. We also learned that uncertainty in building simulations is typically accounted for by simplifying the uncertain parameters to an archetypal probability distribution of possible values.

The field of model predictive control in building systems is also growing to maturity, with numerous examples of simulation studies and a handful of physical implementations occurring in the last decade. The main trends in model predictive control are to employ simplified models for real-time implementation, or complex models for simulation studies, such as those presented in this work.

In the first of four simulation studies presented in this work, the impact of occupant behavior on the performance of a prototypical MM building is explored. The results show that occupant use of windows has the largest impact on HVAC energy consumption, but that comfort is minimally impacted since the HVAC system operates independently of occupant window position. A key finding from this study is that the impact of occupant behavior is not constant relative to changes in automatic building controls. In other words, when one changes the automatic controls in a building, the impact of occupant behavior will also change.

In Chapters 7 and 8, bottom-up and top-down approaches are employed to find how the natural ventilation control logic of the RSF building could be improved. The bottom-up approach showed that incorporating a window-HVAC interlock system was the surest means of capturing energy savings from natural ventilation. Since the ventilation system currently operates independently of the RSF's building-controlled windows, fans continue to run even when natural ventilation is in operation; savings should be easily attainable by enabling fans to ramp down or turn off when natural ventilation is in effect. Without an interlock system, adjustments of a few key setpoint values showed potential for roughly 5% HVAC energy savings on an annual basis. If, instead of annual setpoint values, monthly values are optimized, savings estimates increase to roughly 7.5%, and if improved setpoints are combined with an interlock system, the savings increase to 15%. Noting that these estimates are for a building that is already operating **extremely** efficiently, it stands to reason that savings might be greater for more typical buildings.

Results of the top-down approach provide insight into the differences between stochastic and deterministic MPC through an optimization study of the RSF automatic window controls. Deterministic results tend to over-estimate the savings potential of MPC, and only provide a point estimate of savings. Stochastic results provide a range of savings estimates, and a sequence of control actions that is more robust to the impact of occupant behavior than default controls.

In an attempt to determine the transferability of extracted control rules across different seasons and climates, Chapter 9 uses results of multiple offline MPC investigations to show that rules tend to be location and time specific, and are not transferable in general. The first conclusion is that rules extracted from single-month MPC results are very specific to each month, however rules extracted from data for single months that bridge the cooling and swing seasons (June and September) tend to perform better across seasons than rules extracted from data from the heart of either season. A rule trained on data combined from seven individual one-month results does, on average, perform better than any rule trained on data from a single month alone, when both comfort and HVAC energy consumption are considered.

Results of the cross-climate rule extraction study showed results similar to the cross-season study; rules trained on data from one location typically do not perform well when tested in other locations. Just as certain one-month datasets yielded control rules that were somewhat transferable, data from MPC runs for certain locations lead to control rules that do perform well in multiple other locations.

11.2 Pragmatism

Practically speaking, what can this research bring to the real world? Consider that the largest barriers to completing this project were in modeling and computing. Modeling occupant behavior, modeling the buildings, and integrating the models together are tedious, difficult, time consuming tasks. The MPC environment was also a large investment in time, but really only needs to be done once, and can be used with any number of buildings. Each building, however, will need its own individual model, and it remains to be seen whether occupant behavior models are truly applicable to multiple different buildings, with different occupants, in different places.

With modeling tools improving, and new automated modeling techniques entering the market, the cost of creating detailed building energy models will fall, and one of the major hurdles will diminish in size. Computing costs are also coming down with the advent of widespread and cheaply available cloud computing. With those two major obstacles either gone or considerably smaller, work like the research presented here will be much easier to conduct, and controls optimizations could become a routine part of any building design or retrofit process.

Now consider that the cost of energy is constantly climbing, and that weather forecasts are always improving. The former will make the incentive to save energy grow, and the latter will make it easier to plan ahead for future weather and its impact on building performance. Current day-ahead weather forecasts are highly reliable, and given that the dynamics of most buildings only matter on a time scale of hours to a day or two at most, this amount of future information is adequate for most buildings to plan ahead and schedule operation to combat the weather.

All of this is to say that predictive control is coming to buildings. The real question then becomes: what is our objective, and how do we properly ask our control system to achieve that objective? Throughout this work two objective functions were used to guide optimizers in adjusting building controls in order to save energy and satisfy comfort requirements. Current control systems generally only work towards achieving a single setpoint, but in the future there will be controls designed to maintain indoor environmental quality while limiting energy consumption, or limiting peak demand, or limiting equipment cycling, or some other goal. In many cases, the objectives of a control system may be conflicting, so finding a proper method for checking and balancing multiple objectives is key.

With smartphones in most every person's pocket these days, it is feasible to know where every person (with a smartphone) is at all times, which could enable building systems to anticipate the arrival of occupants, and know exactly where they are at all times. The challenges in this scenario are in implementation and security, but the benefits of a well conditioned building, and a building that is conditioned only when needed may well justify the investment in figuring out how to securely implement such a system.

Occupant Behavior

In Chapter 6, we showed through a simulation study that it is impossible to know a priori what the impact of occupant behavior will be on building performance. We can predict with some accuracy what the impacts will be, but not without an extensive simulation study, which comes after all of the model development and integration work. Without the aid of such a study, what is the best way to account for occupant behavior? Bring occupants into the conversation. Educate occupants on what they should and should not do, tell them what will work and will not work in terms of adjusting their environment, and why. If a building is designed to work with occupants, e.g. through manual windows and night ventilation, then it is critical that occupants know when and why they should open and close windows. This knowledge could come from some rules of thumb that are passed on to occupants through organized instruction, or through a real-time notification system, as in the RSF building.

In the RSF Building, occupants have access to a reporting system, through which they can log their comfort, discomfort, or comments about events and conditions in the building, and through which they receive instructions on when it is a good idea to open or close windows. The ability to quickly and easily log their discomfort empowers occupants, giving them a stronger sense of ownership of and connection to their space, while at the same time alerting facilities management to potential faults in the HVAC system, or in the building design. The ability to ask occupants to close and open windows on command can also eliminate a great deal of uncertainty around when and why they might do so.

We have to be careful though, not to go too far in telling occupants exactly what to do to change their environment. A great deal of occupant satisfaction comes from being able to make adjustments according to one's own internal comfort thresholds and preferences. Indeed, since there is no one-size-fits-all temperature, airflow, or lighting level, we can probably come much closer to optimal by getting close on a large-scale, and allowing occupants to adjust the last bit of lighting (task light), airflow (individual diffuser), or temperature (local thermostat).

Bottom Up

The bottom-up method demonstrated in Chapter 7 is very straightforward in both designing a study and in analyzing results. Such simulations studies should be a part of every building design and retrofit process, assuming a building energy model already exists. Existing design work generally involves a limited simulation study that might investigate variations on window-to-wall ratio, insulation levels, glazing types, different HVAC systems, and other physical properties. The construction and components in a building are only a part of the whole, however, and it is proper operation of these systems that will lead to consistent and more accurate predictions of performance. Controls are often overlooked in the design of a building, but can and do play a very important role in its day to day performance, this is especially true for newer, more complex, high performance buildings that often contain a multitude of systems that have to work in concert to work well.

Top Down

The top-down method shown in Chapters 8 and 9, is considerably more involved, and relies on a broad set of knowledge including building physics, modeling, controls, optimization, statistical modeling, and an ability to work with a considerable amount of data. While results for the RSF building showed minimal potential improvements, the fact that results were positive at all for such an efficient building lends some credibility to the method. In buildings with complex, nonlinear, interacting systems, such methods may go a long way towards finding simple, effective control strategies, and can certainly shed light on the potential upper bound for a system's performance.

One major warning regarding MPC and rule extraction for researchers and practitioners alike is to carefully consider the objective function. Every decision the optimizer makes is guided by the objective function it is given, and if there is some way to exploit a flaw in the building model to improve the value of the objective function, the optimizer will find it, and exploit it. For example if heat is not included in the objective function, the optimizer can find massive cooling savings from natural cooling, but spend massive amounts on re-heating cool air as it enters the building.

Recommendations

As a building designer, what might one do to mitigate the impact of occupant behavior? Limit an occupant's potential to impact building performance by constraining the size of any system they can adjust. If there is only one light switch for an entire wing of a building, and a single occupant is working late, he or she will have to turn on the lights for the entire wing. If, however, the same occupant has a task light or a switch controlling a smaller area of overhead lights, the impact of that occupant working late and keeping the lights on will be smaller. The same logic works for a thermostat, or ventilation system; localized control will give each occupant the ability to impact a smaller portion of the building, without having a large impact on the whole. Give occupants what they need, and not more than they need. The only other major recommendation is to implement a reset system, so that if and when occupants change a system and leave it in an undesirable setting (e.g. window left open on a hot day, blind left down even though outdoor lighting is good), it can be reset to a desirable setting. This could be windows that close automatically at certain times, or thermostats that reset at a reasonable interval.

Finally, ensure that systems operate as intended. If a system is designed to save energy by turning off at night, take measures to ensure that it does in fact turn off at night. Take the RSF building as an example again, it was designed to save cooling energy using night cooling from natural ventilation - but does not, because the automatic windows are not allowed to open when it is cool outdoors. Additionally, the natural ventilation system is intended to offset mechanical ventilation and save energy by turning fans down - but does not, because the mechanical ventilation system operates independently of the natural ventilation.

11.3 Major Contributions

Recall that the main goal of this project was to develop control strategies for conserving energy in a mixed mode building while respecting occupant comfort and accounting for occupant behavior. The chosen building for testing, the RSF building, has set the bar for efficient design, construction, and operation of a large office building, and as such made the goal of the project that much harder to achieve. The RSF is and was already operating very efficiently, and the massive radiant conditioning systems in the building are not very susceptible to occupants opening manual windows, which was the primary occupant behavior considered. Thus, results of the many simulation and optimization studies show only minor potential improvements in building performance, however several important lessons and techniques were learned along the way.

In the investigation of occupant behavior on a prototypical MM building in Chapter 6, we learned that the impact of occupant behavior changes whenever the operation of the building systems change. That is to say, there is a two-way interaction between building systems and manually operable systems which is difficult (if not impossible) to determine a priori without a large simulation study. This is a new finding, and should be considered in future simulation or optimization studies where occupant behavior plays a large role.

Most simulation studies that include stochastic effects use predefined distributions of each stochastic variable, and use guided sampling techniques to efficiently compute the impact of stochastic disturbances on the system being simulated. In this work we used unadulterated and unsimplified stochastic models in a Monte Carlo analysis to learn what the true distribution of results is. The inclusion of occupant behavior models and MC sampling within a building controls optimization problem was tried for the first time in this work. The augmentation of the MPC software to deal with stochastic models and MC sampling, as well as the methodology for treating inputs and outputs during each step of the simulation and optimization process are both contributions that researchers can leverage in future work. The chapter on Control Approaches (Chapter 4 can serve as a how-to guide for future researchers.

Finally, the large simulation, optimization, and rule extraction studies presented in Chapter 9, have shown that extracted control rules for buildings are not transferable, as was hoped. The good news for practitioners is that every building and every climate will require unique control logic, there is no silver-bullet control strategy that will work in all situations.

11.4 Conclusions

The major insights gleaned from the presented work focused on mixed mode buildings, occupant behavior, model predictive control, and the worlds largest net-zero energy building are as follows. First, the impact of occupant behavior on mixed mode buildings can be significant, as the literature shows, and as the global building stock trends towards more efficient designs that include natural ventilation and manually operable systems, the net impact of occupant behavior in buildings will grow. In studies focused on the RSF building, the impact of occupant behavior is less significant since the RSF is designed with systems to operate independently of occupant interaction. Other buildings, i.e. those that rely on occupant window use for adequate ventilation or for night time cooling will be more susceptible to occupant behavior.

Accounting for occupant behavior via the Monte Carlo simulation strategy employed in this work is certainly one of the most accurate methods, but requires considerable work on the front end in implementation and on the back end for analysis. For strategies like this to gain traction on any scale, the barriers to implementation and use of behavioral models need to be removed. The means to easily implement occupant behavior models into building energy simulation, and for easily analyzing the results of multiple simulations are currently unavailable. Additionally, each model of occupant behavior is at least partially biased towards the pool of occupants or buildings that were used to develop the model, and a standard procedure for adapting each model to a given building or climate is needed. Fortunately, the recently founded IEA Annex 66 project[1] will begin work in exactly these topics starting in November, 2014.

Until mixed mode buildings are more common and design paradigms are established,

[1] http://www.annex66.org/

each MM building will continue to feature its own unique control strategy. While the methods applied in this research showed promise for improving upon conventional natural ventilation control heuristics, they relied upon a detailed and accurate energy model of the building and considerable coding to set up and analyze the results of controls studies. For the RSF building, the added work involved in the top-down approach is not justified, since the building is already operating very efficiently. The bottom-up method did, however, identify several opportunities for savings.

In the end, the bottom-up method showed a potential 7% improvement in performance for the RSF building, while the top-down method showed a potential 10% performance improvement by the optimizer, but no savings were possible with extracted rules. In the end this comes from two sources; the lack of savings from rule extraction is partially an artifact of the limited opportunity for improvement in the RSF building, and partially due to the disconnect between statistical models and physical models. Plenty of extracted rules were found that performed well in a statistical sense compared to their own training data, but performed poorly in closed-loop simulations.

The field of rule extraction is still new in the buildings industry, and there is plenty of room for improvement and testing of different rule extraction methods. The concept of converting state information to state-change information for extracting control rules was new in this work, and was shown to improve the performance of extracted rules in closed-loop simulations, but consistent performance in closed-loop simulations was beyond our reach. This leaves the door open for a more comprehensive investigation of how the rule extraction process can be changed to further improve closed-loop performance.

11.5 Future Work

Unanswered Questions A great deal of time and effort was expended during this project to work with a single building, the RSF, so most of the results can be presented with a footnote that reads “these results are specific to the RSF”. It would be interesting to apply the methodology developed throughout this work to a broad range of buildings, including not just mixed-mode buildings and those with manually operable systems, but also to conventional buildings with conventional controls to see if they can be improved upon, and if so, how. Another obvious next-step is to try the methodology on a building that truly needs significant improvement, to see how much more the top down approach can improve performance than the bottom up (if at all).

It would also be interesting to apply the rule extraction techniques to different control problems for different buildings. How much data does one need for rule extraction to work properly, what are the best inputs to provide to the CART (or other) machine learning algorithm? What are the best set of CART growth and pruning parameters? In this work we followed rules of thumb and used default parameters for guiding the CART algorithm, but there are numerous knobs one can turn to get different results from the CART algorithm. In addition to digging deep into the CART algorithm to see how it could be tuned, another goal is to investigate a set of rule extraction techniques, including random forests, adaptive boosting, and others, to see what machine learning algorithms worked best in building controls.

Occupant Behavior Model Comparison In the work that was restricted to investigating the impact of occupant behavior, only four different models were considered, one each for lighting, shading, windows, and occupancy. The literature lacks any comprehensive comparison of numerous different models, though several models exist for predicting occupancy, and for occupant use of windows. While verbose comparisons exist, there is no simulation study that would compare the models on even terms, with the same building model, for example.

Cost Functions A crucial part of every MPC investigation that can not be underscored enough is the formulation of the cost function. In this work we considered energy consumption and occupant comfort, as defined by the Fanger PMV/PPD model. In many MPC in buildings studies, comfort is considered by establishing limits on indoor air temperature, and the optimizer is asked to keep temperatures within the limits. In our work, we asked the optimizer to reduce energy consumption, and improve comfort, but in reality there will always be tradeoffs, and establishing a sound method for fairly trading energy and comfort could change the cost function, which could in turn change the optimization results considerably. For future work, a comparison of the cost functions used in current MPC in buildings investigations, and an attempt to establish what works best are needed. Currently every MPC investigation uses its own cost function, just as every MM building uses its own new set of controls.

MPC or Rule Extraction? Given the heavy computational burden and scientific overhead in the top-down approach, including rule extraction, one might ask: why bother? Why not just embed a building's controller with the MPC algorithm and control it directly? The answer is, and will continue to be, a lack of faith from the controls contractors and building operators that care for our buildings every day. The HVAC industry is slow to adapt to change and to adopt new technologies, and will continue to reject the black-box that is an MPC controller for years to come. Rule extraction provides a compromise between next-generation control and current control practice. Control rules that are readable, consistent, and transparent. Even to an expert an MPC controller may appear to be misbehaving at times due to some nuance of the objective function or the forecasted model that a casual observer cannot see. In the end, all control strategies are born out of experience and testing; in the case of rule extraction from optimal data sets, the control strategies are the child of an optimizer's testing and experience with a model.

Bibliography

[1] René Altherr and Jean-Bernard Gay. A low environmental impact anidolic facade. Building and Environment, 37(12):1409–1419, December 2002. 00031.

[2] Rune Vinther Andersen, Bjarne W. Olesen, and Jørn Toftum. Simulation of the effects of occupant behaviour on indoor climate and energy consumption. In Clima 2007 WellBeing Indoors, Halsinki, Finland, 2007.

[3] ASHRAE. Energy standard for buildings except low-rise residential buildings, 2004.

[4] ASHRAE. Thermal environmental conditions for human occupancy, 2010.

[5] ASHRAE. Ventilation for acceptable indoor air quality, 2010.

[6] Mustafa Y. Ata. A convergence criterion for the monte carlo estimates. Simulation Modelling Practice and Theory, 15(3):237–246, March 2007.

[7] Nick Baker and Mark Standeven. Thermal comfort for free-running buildings. Energy and Buildings, 23(3):175–182, March 1996.

[8] Alberto Bemporad and Manfred Morari. Robust model predictive control: A survey. Robustness in identification and control, page 207–226, 1999.

[9] Dimitris Bertsimas and Melvyn Sim. Tractable approximations to robust conic optimization problems. MATH. PROGRAM, 107:2006, 2006.

[10] William Bordass, Adrian Leaman, and S Willis. Control strategies for building services: the role of the user. BRE, May 1994.

[11] Samuel Borgeson and Gail S. Brager. Occupant control of windows: Accounting for human behavior in building simulation. University of California, Berkeley, 2008.

[12] Denis Bourgeois, Christoph Reinhart, and Iain Macdonald. Adding advanced behavioural models in whole building energy simulation: A study on the total energy impact of manual and automated lighting control. Energy and Buildings, 38(7):814–823, July 2006.

[13] Gail S. Brager, Samuel Borgeson, and Yoon Soo Lee. Control strategies for mixed-mode buildings. Technical report, Center for the Built Environment, 2008.

[14] James E. Braun. Reducing energy costs and peak electrical demand through optimal control of building thermal storage. ASHRAE transactions, 96(2):876–888, 1990.

[15] Mark Cannon, Basil Kouvaritakis, and Xingjian Wu. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. IEEE Transactions on Automatic Control, 54(7):1626–1632, July 2009.

[16] CEN. Ventilation for non-residential buildings - performance requirements for ventilation and room-conditioning systems, 2007.

[17] Chartered Institution of Building Services Engineers. Mixed mode ventilation: CIBSE applications manual AM 13: 2000. Chartered Institution of Building Services Engineers, London, 2000.

[18] Debasish Chatterjee, Peter Hokayem, and John Lygeros. Stochastic model predictive control with bounded control inputs: a vector space approach. 0903.5444, March 2009.

[19] Darcie Chinnis and Gregor P. Henze. A comparison of lighting energy modeling methods to simulate annual energy use and peak demand. LEUKOS, 9(2):109–126, 2012.

[20] CIBSE. Natural Ventilation in Non-Domestic Buildings. CIBSE Applications Manual 10. Chartered Institution Of Building Services Engineers, 1997.

[21] J. A. Clarke, I. Macdonald, and J. F. Nicol. Predicting adaptive responses - simulating occupied environments. In Proceedings of international comfort and energy use in buildings conference. Network for Comfort and Energy Use in Buildings (NCEUB), London, 2006. 00013.

[22] Brian Coffey. Using Building Simulation and Optimization to Calculate Lookup Tables for Control. PhD thesis, University of California, Berkeley, January 2012.

[23] Raymond J. Cole and Zosia Brown. Reconciling human and automated intelligence in the provision of occupant comfort. Intelligent Buildings International, 1(1):39–55, 2009.

[24] Charles D. Corbin, Gregor P. Henze, and Peter May-Ostendorp. A model predictive control optimization environment for real-time commercial building application. Journal of Building Performance Simulation, 6(3):159–174, 2013.

[25] Richard J. de Dear and Gail S. Brager. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE standard 55. Energy and Buildings, 34(6):549–561, July 2002.

[26] Alexander Domahidi, Fabian Ullmann, Manfred Morari, and Colin N. Jones. Learning near-optimal decision rules for energy efficient building control. In Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, pages 7571–7576, Maui, HI, 2012.

[27] Dept. of Energy. Total energy - data - U.S. energy information administration (EIA). Technical report, US Energy Information Administration, 2010.

[28] Valentina Fabi, Rune Vinther Andersen, Stefano Corgnati, and Bjarne W. Olesen. Occupants' window opening behaviour: A literature review of factors influencing occupant behaviour and models. Building and Environment, 58:188–198, December 2012.

[29] Rolf Findeisen and Frank Allgöwer. An introduction to nonlinear model predictive control. In 21st Benelux Meeting on Systems and Control, volume 11, 2002.

[30] Roberto Z. Freire, Gustavo H.C. Oliveira, and Nathan Mendes. Predictive controllers for thermal comfort optimization and energy savings. Energy and Buildings, 40(7):1353–1365, 2008.

[31] R. Fritsch, A. Kohler, M. Nygård-Ferguson, and Jean-Louis Scartezzini. A stochastic model of user behaviour regarding ventilation. Building and Environment, 25(2):173–181, 1990.

[32] J. M. Fürbringer and C. A. Roulet. Comparison and combination of factorial and monte-carlo design in sensitivity analysis. Building and Environment, 30(4):505–519, October 1995.

[33] Zachary M. Gill, Michael J. Tierney, Ian M. Pegg, and Neil Allan. Low-energy dwellings: the contribution of behaviours to actual performance. Building Research & Information, 38(5):491–508, 2010.

[34] Lixing Gu. Airflow network modeling in EnergyPlus. In Building Simulation, pages 1041–1047, Beijing, China, September 2007.

[35] Markus Gwerder and Jürg Tödtli. Predictive control for integrated room automation. In 8th REHVA World Congress Clima, 2005.

[36] D Haigh. User response in environmental control. In The Architecture of energy, pages 45–63. Longmans ; Construction Press, New York : Burnt Mill, Essex, 1981.

[37] Frédéric Haldi. Towards a Unified Model of Occupants' Behaviour and Comfort for Building Energy Simulation. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2010.

[38] Frédéric Haldi and Darren Robinson. A comprehensive stochastic model of window usage: theory and validation. In Proceedings of Building Simulation, 2009.

[39] Frédéric Haldi and Darren Robinson. Interactions with window openings by office occupants. Building and Environment, 44(12):2378–2395, December 2009.

[40] Frédéric Haldi and Darren Robinson. Adaptive actions on shading devices in response to local visual stimuli. Journal of Building Performance Simulation, 3(2):135–153, 2010.

[41] Frédéric Haldi and Darren Robinson. The impact of occupants' behaviour on building energy demand. Journal of Building Performance Simulation, 4(4):323–338, December 2011.

[42] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. The Elements of Statistical Learning. Springer, corrected edition, July 2003.

[43] James Hauswirth and Gregor P. Henze. Evaluating the tradeoffs of occupant comfort and energy savings: A study of window control sensitivity. In AEI 2013, pages 867–876. American Society of Civil Engineers, 2013.

[44] Philip Haves and Peng Xu. The building controls virtual test bed a simulation environment for developing and testing control algorithms, strategies and systems. In Proc. of the 10-th IBPSA Conference, page 1440–1446, 2007.

[45] Per Kvols Heiselberg. Control strategies for hybrid ventilation in new and retrofitted office and education buildings. Technical report, International Energy Agency, 2006.

[46] Achim Hekler, Daniel Lyons, Benjamin Noack, and Uwe D. Hanebeck. Nonlinear model predictive control considering stochastic and systematic uncertainties with sets of densities. In Proceedings of the IEEE Multi-Conference on Systems and Control (MSC 2010), Yokohama, Japan, September 2010.

[47] Gregor P. Henze. Impact of real-time pricing rate uncertainty on the annual performance of cool storage systems. Energy and Buildings, 35(3):313–325, March 2003.

[48] Gregor P. Henze, Robert H. Dodier, and Moncef Krarti. Development of a predictive optimal controller for thermal energy storage systems. HVAC&R Research, 3(3):233–264, 1997.

[49] Gregor P. Henze, Clemens Felsmann, and Gottfried Knabe. Evaluation of optimal control for active and passive building thermal storage. International Journal of Thermal Sciences, 43(2):173–183, February 2004.

[50] Sebastian Herkel, Ulla Knapp, and Jens Pfafferott. Towards a model of user behaviour regarding the manual control of windows in office buildings. Building and Environment, 43(4):588–600, 2008.

[51] Michael A Humphreys. Understanding the adaptive approach to thermal comfort. ASHRAE Transactions, 104(1):991–1004, 1998.

[52] D.R.G. Hunt. Predicting artificial lighting use - a method based upon observed patterns of behaviour. Lighting Research and Technology, 12(1):7 –14, March 1980.

[53] Dirk Jacob. Optimization techniques for building simulations with uncertain parameters. In 8th International Conference on System Simulation in Buildings, Liège, December 2010.

[54] Ali A. Jalali and Vahid Nadimi. A survey on robust model predictive control from 1999-2006. In Computational Intelligence for Modelling, Control and Automation, 2006 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce, International Conference on, page 207–207, 2006.

[55] James Kennedy and Russell C. Eberhart. A discrete binary version of the particle swarm algorithm. In Systems, Man, and Cybernetics, 1997. Computational Cybernetics and Simulation., 1997 IEEE International Conference on, volume 5, pages 4104–4108, Orlando, FL, USA, 1997.

[56] Adrian Leaman and Bill Bordass. Productivity in buildings: the 'killer' variables. Building Research & Information, 27(1):4–19, 1999.

[57] Kevin J. Lomas and Herbert Eppel. Sensitivity analysis techniques for building thermal simulation programs. Energy and Buildings, 19(1):21–44, 1992.

[58] Peter Jacobus Lute. The use of predictions in temperature control in buildings: A passive climate system application. PhD thesis, Delft University of Technology, October 1992. Prof. Ir. J. Bekker.

[59] Yudong Ma, Francesco Borrelli, Brandon Hencey, Brian Coffey, Sorin Bengea, and Philip Haves. Model predictive control for the operation of building cooling systems. In American Control Conference (ACC), 2010, pages 5106–5111, Baltimore, MD, USA, June 2010. IEEE.

[60] Yudong Ma, Francesco Borrelli, Brandon Hencey, Andrew Packard, and Scott A. Bortoff. Model predictive control of thermal energy storage in building cooling systems. In Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009, pages 392–397. IEEE, December 2009.

[61] Iain Alexander Macdonald. Quantifying the Effects of Uncertainty in Building Simulation. PhD thesis, University of Strathclyde, July 2002.

[62] Alie El-Din Mady, Gregory M. Provan, Conor Ryan, and Kenneth N. Brown. Stochastic model predictive controller for the integration of building use and temperature regulation. In Twenty-Fifth AAAI Conference on Artificial Intelligence, San Francisco, CA, USA, 2011.

[63] Peter May-Ostendorp, Gregor P. Henze, Charles D. Corbin, Balaji Rajagopalan, and Clemens Felsmann. Model-predictive control of mixed-mode buildings with rule extraction. Building and Environment, 46(2):428–437, February 2011.

[64] Peter May-Ostendorp, Gregor P. Henze, Balaji Rajagopalan, and Doreen Kalz. Experimental investigation of model predictive control-based rules for a radiantly cooled office. HVAC&R Research, 19(5):602–615, 2013.

[65] Seung-Ki Min, Xuebin Zhang, Francis W. Zwiers, and Gabriele C. Hegerl. Human contribution to more-intense precipitation extremes. Nature, 470(7334):378–381, February 2011.

[66] National Renewable Energy Laboratory. NREL: sustainable NREL - research support facility, October 2012.

[67] Guy Newsham, Ardeshir Mahdavi, and Ian Beausoleil-Morrison. Lightswitch: A stochastic model for predicting office lighting energy consumption. 1995.

[68] Fergus Nicol and Michael A. Humphreys. Derivation of the adaptive equations for thermal comfort in free-running buildings in european standard EN15251. Building and Environment, 45(1):11–17, January 2010.

[69] J Fergus Nicol and Michael A Humphreys. A stochastic approach to thermal comfort-occupant behavior and energy use in Buildings/Discussion. ASHRAE Transactions, 110:554, 2004.

[70] Kristian Nolde, Markus Uhr, and Manfred Morari. Medium term scheduling of a hydro-thermal system using stochastic model predictive control. Automatica, 44(6):1585–1594, June 2008.

[71] Frauke Oldewurtel, Colin N. Jones, and Manfred Morari. A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback. In Decision and Control, 2008. CDC 2008. 47th IEEE Conference on, pages 4731–4736, Cancun, Mexico, 2008.

[72] Frauke Oldewurtel, Alessandra Parisio, Colin N. Jones, Manfred Morari, Dimitrios Gyalistras, Markus Gwerder, Vanessa Stauch, Beat Lehmann, and Katharina Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In American Control Conference (ACC), 2010, pages 5100–5105, Baltimore, MD, USA, 2010.

[73] Nigel A. Oseland. Predicted and reported thermal sensation in climate chambers, offices and homes. Energy and Buildings, 23(2):105–115, December 1995.

[74] Jessen Page. Simulating Occupant Presence and Behaviour in Buildings. PhD dissertation, Ecole Polytechnique Federale de Lausanne, 2007.

[75] Jessen Page, Darren Robinson, Nicolas Morel, and Jean-Louis Scartezzini. A generalised stochastic model for the simulation of occupant presence. Energy and Buildings, 40(2):83–98, 2008.

[76] Pardeep Pall, Tolu Aina, Dáithí A. Stone, Peter A. Stott, Toru Nozawa, Arno G. J. Hilberts, Dag Lohmann, and Myles R. Allen. Anthropogenic greenhouse gas contribution to flood risk in england and wales in autumn 2000. Nature, 470(7334):382–385, February 2011.

[77] Wout Parys, Hilde Breesch, Hugo Hens, and Dirk Saelens. Feasibility assessment of passive cooling for office buildings in a temperate climate through uncertainty analysis. Building and Environment, 56:95–107, October 2012.

[78] Wout Parys, Dirk Saelens, and Hugo Hens. Coupling of dynamic building simulation with stochastic modelling of occupant behaviour in offices – a review-based integrated methodology. Journal of Building Performance Simulation, 4(4):339–358, December 2011.

[79] Scott Pigg, Mark Eilers, and John Reed. Behavioral aspects of lighting and occupancy sensors in private offices: A case study of a university office building. ACEEE 1996 Summer Study on Energy Efficiency in Buildings, 1996.

[80] ASHRAE Press. The ASHRAE GreenGuide, Second Edition. Butterworth-Heinemann, 2 edition, November 2006.

[81] Christoph F. Reinhart. Lightswitch-2002: a model for manual and automated control of electric lighting and blinds. Solar Energy, 77(1):15–28, 2004.

[82] Hom B. Rijal, Michael A. Humphreys, and J. Fergus Nicol. Understanding occupant behaviour: The use of controls in mixed-mode office buildings. Building Research and Information, 37(4):381–396, 2009. Compendex.

[83] Hom B. Rijal, Paul Tuohy, Michael A. Humphreys, J. Fergus Nicol, and Aizaz Samuel. An algorithm to represent occupant use of windows and fans including situation-specific motivations and constraints. Building Simulation, 4(2):117–134, December 2011.

[84] Hom B. Rijal, Paul G. Tuohy, Michael A. Humphreys, J Fergus Nicol, Aizaz Samuel, and Joseph Clarke. Using results from field surveys to predict the effect of open windows on thermal comfort and energy use in buildings. Energy and Buildings, 39(7):823–836, July 2007.

[85] Horn B. Rijal, Aizaz Samuel, Paul Tuohy, Michael A. Humphreys, Iftikhar A. Raja, Joe Clarke, and J. Fergus Nicol. Development of adaptive algorithms for the operation of windows, fans, and doors to predict thermal comfort and energy use in pakistani buildings. ASHRAE Transactions, 114(2):555–573, October 2008.

[86] Darren Robinson and Frédéric Haldi. Modelling occupants' presence and behaviour – part i. Journal of Building Performance Simulation, 4(4):301–302, 2011.

[87] Astrid Roetzel, Aris Tsangrassoulis, Udo Dietrich, and Sabine Busching. A review of occupant control on natural ventilation. Renewable and Sustainable Energy Reviews, 14(3):1001–1013, April 2010.

[88] Claude-Alain Roulet, Niklaus Johner, Flavio Foradini, Philomena Bluyssen, Chrit Cox, Eduardo De Oliveira Fernandes, Birgit Müller, and Claire Aizlewood. Perceived health and comfort in relation to energy use and building characteristics. Building Research & Information, 34(5):467–474, 2006.

[89] O Seppänen and W J Fisk. Association of ventilation system type with SBS symptoms in office workers. Indoor air, 12(2):98–112, June 2002. 00151 PMID: 12216473.

[90] Jan Široký, Frauke Oldewurtel, Jiří Cigler, and Samuel Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. Applied Energy, 88(9):3079–3087, September 2011.

[91] Richard K. Strand, Drury B. Crawley, Curtis O. Pedersen, Richard J. Liesen, Linda K. Lawrie, F. C. Winkelmann, W. F. Buhl, J. Huang, and Daniel E. Fisher. EnergyPlus: a new-generation energy analysis and load calculation engine for building design. In Proceedings of the 2000 ASCA Technology Conference, Cambridge, MA, USA, 2000.

[92] Ryan A. Tanner and Gregor P. Henze. Quantifying the impact of occupant behavior in mixed mode buildings. In AEI 2013, pages 246–255. American Society of Civil Engineers, 2013.

[93] Ryan A. Tanner, Gregor P. Henze, and Shanti Pless. Optimized control of automatic windows for energy savings and occupant comfort (1597-RP). ASHRAE Transactions, 2013.

[94] Paul G. Tuohy, Michael A. Humphreys, J. Fergus Nicol, Hom B. Rijal, and Joe A. Clarke. Occupant behavior in naturally ventilated and hybrid buildings. ASHRAE Transactions, 115(1):16–27, May 2009.

[95] Energy Information Administration (US). International Energy Outlook 2013 With Projections to 2040. Government Printing Office, 2013.

[96] EIA US Energy Information Agency. Annual energy review 2011. Technical report, US Energy Information Agency, September 2012.

[97] Christine E. Walker, Leslie K. Norford, and Leon R. Glicksman. Assessing the performance of a naturally ventilated office building. In ASHRAE transactions, volume 112, pages 281–294. American Society of Heating, Refrigerating and Air-conditioning Engineers, 2006.

[98] Danni Wang, Clifford C. Federspiel, and Francis Rubinstein. Modeling occupancy in single person offices. Energy and Buildings, 37(2):121–126, February 2005.

[99] P. R. Warren and L. M. Parkins. Window-opening behaviour in office buildings. Building Services Engineering Research and Technology, 5(3):89–101, January 1984.

[100] Michael Wetter. Co-simulation of building energy and control systems with the building controls virtual test bed. Journal of Building Performance Simulation, 4(3):185–203, November 2011. 00083.

[101] Jun Xu, Peter B. Luh, William E. Blankson, Ron Jerdonek, and Khalil Shaikh. An optimization-based approach for facility energy management with uncertainties. HVAC&R Research, 11(2), 2005.

[102] Yohei Yamaguchi, Yoshiyuki Shimoda, and Minoru Mizuno. Development of district energy system simulation model based on detailed energy demand model. In Proceeding of Eighth International IBPSA Conference, page 1443–1450, 2003.

[103] Geun Young Yun and Koen Steemers. Time-dependent occupant behaviour models of window control in summer. Building and Environment, 43(9):1471–1482, September 2008.

[104] Geun Young Yun, Paul Tuohy, and Koen Steemers. Thermal performance of a naturally ventilated building using a combined algorithm of probabilistic occupant behaviour and deterministic heat and mass balance models. Energy and Buildings, 41(5):489–499, May 2009.