Model Guided Control of Radiant Slabs for Comfort and Efficiency

A Dissertation

Presented in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

with a

Major in Mechanical Engineering

in the

College of Graduate Studies

University of Idaho

by

L. Damon Woods

Co-Major Professor: Ralph Budwig, Ph.D.

Co-Major Professor: Kevin Van Den Wymelenberg Ph.D.

Committee Members: John Gardner, Ph.D.; Behnaz Rezaie Ph.D.

Department Administrator: Steven Beyerlein, Ph.D.

August 2018

AUTHORIZATION TO SUBMIT DISSERTATION

This dissertation of Damon Woods, submitted for the degree of Doctor of Philosophy with a Major in Mechanical Engineering and titled “Model Guided Control of Radiant Slabs for Comfort and Efficiency," has been reviewed in final form. Permission, as indicated by the signatures and dates given below, is now granted to submit final copies to the College of Graduate Studies for approval.

Co-Major
Professors: ______________________________ Date: __________
Ralph Budwig, Ph.D.

______________________________ Date: __________
Kevin Van Den Wymelenberg, Ph.D.

Committee
Members:
______________________________ Date: __________
John Gardner, Ph.D.

______________________________ Date: __________
Behnaz Rezaie, Ph.D.

Department
Administrator:
______________________________ Date: __________
Steven Beyerlein, Ph.D.

ABSTRACT

Radiant slabs have the potential to deliver efficient heating and cooling to buildings while enhancing comfort. However, these systems do not operate on conventional time scales and there is often a delay between a thermostat’s signal and the system’s response. This research lays out a framework of how to overcome this problem by linking an energy model to weather forecasts to inform the controls. An overview of radiant systems and technologies demonstrates their inherent advantages for conditioning buildings. A method for connecting an energy model to control hardware for remote co-simulation was developed and tested. Results indicated operational savings of at least 12% could be achieved through the remote co-simulation approach. Once it was verified that an OpenStudio model could be used for virtual commissioning, the focus turned towards modeling of a radiant slab. Extensive data was collected at a design office in Boise, Idaho with a radiant slab used for heating and cooling. An OpenStudio model was built and calibrated to predict the heating and cooling required of the radiant slab in real time. A framework was set up for performing parametric simulations and incorporating weather forecasts into the energy model. The model was used to predict ideal control setpoints by using short simulations of four hours into the future. Model comparisons indicated that using the model to guide the control setpoints could save up to 13% of HVAC energy over one month and the office space was kept significantly more comfortable compared to the current control scheme.

ACKNOWLEDGEMENTS

I would like to thank Dr. Ralph Budwig who has been a tireless advocate of my work and provided encouragement throughout this project. Your advisement has been invaluable. Thank you to Dr. Kevin Van Den Wymelenberg for bringing me into the Integrated Design Lab, demonstrating a passion for the work, and pushing my research to its best. Thank you to Dr. Behnaz Rezaie for your advice on the project and publications. Thank you to Dr. John Gardner for inspiring me; you’ve opened my eyes to the world of energy efficiency and given me a clear understanding of what I want to study from now on. This project would be impossible without tremendous support from the wonderful research team at the Integrated Design Lab. IDL has become my second home and I’ve learned so much from everyone there. Elizabeth Cooper, thank you for your mentorship, kindness, and flexibility while I tried to balance work and school. Funding for this project has come in various parts through generous support from the Northwest Energy Efficiency Alliance, the Center for Advanced Energy Studies, Idaho Power, and Avista. Thank you Russ Pratt and Kent Hanway at CSHQA for providing access to your building and answering my many questions. Thank you Sean Rocke, Jake MacArthur, and Sean Rosin for helping me navigate EMS systems.

DEDICATION

For my parents

for their ever-present encouragement and inspiring example.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

A/D	Analog Digital
AHU	Air Handling Unit
AMY	Actual Meteorological Year
API	Application Programming Interface
ASHRAE	American Society of Heating Refrigeration and Air Conditioning Engineers
BACnet	Building Automation and Control network
BCVTB	Building Controls Virtual Test Bed
CBE	Center for the Built Environment
CFM	Cubic Feet per Minute
DOAS	Dedicated Outdoor Air Systems
DOE	U.S. Department of Energy
EEM	Energy Efficiency Measure
EIA	Energy Information Administration
EMS	Energy Management System
EPW	EnergyPlus Weather File Format
GPM	Gallons Per Minute
HP	Horespower

HVAC	Heating Ventilating and Cooling
IAQ	Indoor Air Quality
LBNL	Lawrence Berkeley National Laboratory
LEED	Leadership in Energy Efficient Design
MA	Mixed Air
MATLAB	Matrix Laboratory Software
MPC	Model Predictive Control
MSTP	Master Slave Token Passing
NOAA	National Oceanographic and Atmospheric Association
NREL	National Renewable Energy Laboratory
NSRDB	National Solar Resource Database
OA	Outdoor Air
RA	Return Air
ROM	Reduced Order Model
RMSE	Root Mean Square Error
RTU	Rooftop Unit
STP	Standard Temperature and Pressure
TMY	Typical Meteorological Year
VAV	Variable Air Volume

WAHP	Water to Air Heat Pump
WWHP	Water to Water Heat Pump

PREFACE

Several chapters of the dissertation draw upon previously published work, some of which has been submitted to the peer-review process. I, Damon Woods, was either the sole or lead author of the publications included in this work. The research was funded by a variety of sources and would not be possible without it.

Chapter 1 on radiant systems draws heavily from the NEEA radiant design guide. The design guide was funded by the Northwest Energy Efficiency Alliance [1]. I supplied the text and calculations while the formatting and some of the figures were developed by Dylan Agnes at the Integrated Design Lab. The second half of Chapter 1 owes much of the text to “Using the Time Delay of Radiant Systems to the Grid’s Advantage” [2]. The work for this publication was funded by the Center for Advanced Energy Studies Energy Efficiency Research Institute. I was the main author of the text and findings; Dr. John Gardner, co-author, wrote the proposal, and oversaw the research and editing of the paper.

The work in Chapter 2 on advances in building controls includes text and figures from “Simulation on Demand for Deep Energy Retrofits” [3]. My role was to update the energy models to a newer version and help develop some of the measures including some coding in Ruby. Co-authors included Alen Mahic, Kevin Van Den Wymelenberg, John Jennings, and Jeff Cole. Mr. Mahic wrote the bulk of the measures and helped oversee the modeling process along with Mike Hatten. Kevin Van Den Wymelenberg oversaw the modeling team and provided feedback on the publication text. John Jennings and Jeff Cole provided oversight and guidance of the overall project.

Chapter 3 on communication between buildings and energy models is drawn from "Optimizing Economizer Operation by Virtual Commissioning through Remote Co-simulation" [4]. This was funded by Avista. Co-authors included Tyler Noble, Brad Acker, Ralph Budwig, and Kevin Van Den Wymelenberg. My role in the project was to execute the tasks of the HVAC modeling and setting up communication between the model and the controller along with Tyler Noble. I supplied the initial text of the publication. Brad Acker developed and proposed the research idea. Dr. Budwig oversaw Tyler's research portion of the project and provided editing and feedback on the paper. Kevin Van Den Wymelenberg managed the project and also provided feedback on the paper.

Chapter 4 on using forecasts and models to guide the controls was funded in part by Idaho Power Company in research that overlapped with this project. Kevin Van Den Wymelenberg secured the funding for this program. Jacob Dunn proposed the initial monitoring of the site and helped document the operation. Shrief Shrief and Brad Acker assisted in the first phase of data collection. Russ Pratt of CSHQA provided access to the controls and informed the team on the operational setpoints. In the second phase of data collection, Jake MacArthur of BuildingFit helped me to install SkySpark at the site for recording control signals. My work was to analyze the data, build the energy model, build measures to report model outputs and to run the alternative control schemes.

CHAPTER 1: RADIANT SYSTEMS – ADVANTAGES AND DYNAMICS

Introduction

Energy conservation and occupant comfort are often viewed as competing interests in the building sector. Radiant systems have the potential to help resolve this conflict by improving indoor environments while minimizing energy inputs. New thermostat controls are needed for radiant systems to achieve their potential. Linking radiant systems to model guided controls can optimize comfort and efficiency. The research presented here demonstrates how to leverage weather forecasts, live operational data, and standard comfort models for a holistic energy management approach.

This control approach is aimed at combating global warming without inhibiting productivity. While this may not be immediately apparent, consider that this technology is targeting the largest category of use in the highest energy sector. Buildings consume more energy (41%) than either transportation (28%) or industry (31%). Of that building energy, close to half (44%) is consumed by Heating Ventilating and Air Conditioning (HVAC) [5]. In economic terms, Americans spend approximately $175 billion each year for the energy to run HVAC [6]. The greenhouse emissions from air-conditioning alone in the U.S. add up to nearly 500 million tons of CO_2 equivalent per year – greater than that of the entire construction industry [7]. Therefore, any improvements in HVAC efficiency result in enormous upstream savings both economically and environmentally.

Based on the preceding statistics, some might advocate for reducing HVAC consumption by adjusting occupant expectations and behavior. In colloquial terms, this option is presented as: “be uncomfortable for the environment’s sake.” Some leaders still advocate this altruistic approach. Unfortunately, changing occupant behavior can be challenging; importantly, worker productivity drops dramatically when the environment is thermally uncomfortable [8]. The purpose, or telos, of the HVAC systems is to provide comfort and the better they provide comfort, the more virtuous they are. Thus, ideal/virtuous controls will manage for the thermal profile of the space – considering not just the air temperature, but the surface temperatures as well to address comfort more holistically.

The research objective is to explore ways of providing thermal comfort while minimizing the environmental expense by combining the technologies of operative temperature controls, radiant systems, and energy modeling. It is hoped that this integration will spark new conversations and design ideas so that people no longer have to choose between comfort and environmental responsibility. HVAC savings can be found through the adoption of more innovative HVAC controls and a wider recognition of the benefits of radiant systems.

Radiant HVAC systems hold great economic potential because they meet the needs of both building operators and grid operators. Traditionally, these interests are often at odds with one another. The objective of a building operator is to create a comfortable environment and minimize occupant complaints of thermal discomfort, while the objective of a grid operator is to ensure enough resources are available to meet the power demands in a region. In order to maintain occupant comfort, building operators often run the cooling air-conditioning systems at their maximum capacity during the warmest part of the day. When all offices in a city respond to the rising afternoon temperatures in the same way, it results in peak power draws that put intense strain on the electrical grid. Radiant systems by contrast require a very different cooling strategy, which reduces peak consumption and adds welcome diversity to the grid consumption profile.

Radiant HVAC systems have been receiving increased attention in the building community [9]. They are increasingly cost competitive and have been shown to provide better thermal comfort than all-air systems [10]. Radiant systems are able to use far less energy than all-air HVAC systems by virtue of their thermal transport property differences. Radiant systems move thermal energy through water instead of air. Water has a much higher density and thermal capacitance than air, allowing a small quantity of water to be pumped through the building to provide cooling instead of using fans to force a large volume of air through the space. Radiant systems reach their peak loads at different times than conventional all-air systems, shifting the phase in the demand profile [11]. A group of buildings operating with radiant systems can provide a counterbalance to the typical load profiles of buildings with air-based heating and cooling. Because of this, a building operator with a radiant slab can run the system in a way that maximizes occupant comfort without incurring high energy consumption during utility peak hours.

While market interest is growing, there are comparatively few radiant systems being installed today [12]. Legacy radiant systems were often managed poorly, either through miscommunication or a lack of sophisticated controls. Radiant systems are unique in that there is a long delay between when a control signal is given and when the thermostat in that zone senses a change in air temperature. Using traditional control methods with radiant systems can add to the peak load and create occupant discomfort due to the lag between control signals and responses. Conversely, when a well-designed control scheme is implemented that does not rely on thermostatic temperature alone, radiant systems can increase thermal comfort, lower energy consumption, shift electrical loads to off-peak hours, and can act as thermal batteries allowing for demand response capabilities and increased grid resiliency.

The dissertation objective is to reconcile the gap between the control theory and the practice of installing and operating radiant slabs in buildings. Without the proper communication, commissioning, and controls in place, few of the benefits of radiant slabs will be realized. Therefore, communication between industry professionals – operators, design engineers, architects, and contractors – is essential. This information exchange helps identify common barriers to radiant HVAC installation and brings to light areas of improvement. On a local level, this exchange included contact with a specific site – a 20,000 ft^2 single-story office building in Boise, ID. At the state level, this research benefits the regional utility as it works to determine energy efficiency incentives for businesses that choose to install radiant cooling. More broadly, these tools and strategies may be helpful for anyone considering a new radiant system or updating a building’s control settings.

Research Objectives and Methods

The research objective was to create a control strategy for a radiant slab that relied on predictive capabilities to deliver excellent occupant comfort. The research began with a review of radiant systems, their dynamics and capabilities. The focus of the research was one particular case study of an office in Boise Idaho that uses a radiant slab as its primary HVAC system. Having a grasp of the radiant HVAC components allowed insight into how these systems ought to be controlled. The research then turned to different types of predictive controls. An overview is provided of the traditional Model Predictive

Controls that rely on a set of differential equations. However, instead of using equations for the model, a full energy model of the site was built in EnergyPlus to approximate building behavior. A parametric simulation approach was developed so that the energy model could be used as part of a predictive control strategy. This is referred to as Model Guided Control. The next step in the research was to ensure that the energy model could communicate with a building's control system. For this phase of the research, a controller was taken from an operational site and connected to an energy model. The research team used the energy model to determine ideal control setpoints for that building in a virtual commissioning process. Having established a parametric simulation approach, and verified the potential of virtual control commissioning through energy models, the research focus shifted towards predictive capabilities. A method is presented for developing local weather forecasts that can be used in model simulations. Now, the energy model of the radiant office can be run in a predictive mode based on these forecasts. Looking ahead, the energy model can adjust the control step at different intervals with the objective of optimizing comfort during occupied periods, and minimizing energy use during unoccupied periods. This control strategy was contrasted with the current operations and shows the potential of improving comfort and reducing energy use. This method could be applied to any building type but has particular relevance for buildings with radiant systems which can benefit most from predictive controls. Radiant systems are also very efficient and can provide great comfort.

Advantages of Radiant Systems

Radiant systems have become increasingly popular methods of heating and cooling a space for several reasons. By conducting heat in and out of buildings through water instead of air, ductwork can be minimized, smaller equipment used, and occupant comfort increased [13]. A radiant system is any HVAC system that meets more than half of its heating or cooling load through radiant means [14]. For occupied buildings, these systems must be combined with a Dedicated Outdoor Air System (DOAS) to provide fresh air in any occupied space. Radiant systems use 40-58% less energy than conventional all-air systems [15]. They can provide improved thermal comfort while operating with much smaller equipment than conventional systems [16].There are three compelling reasons that may sway a client to adopt a radiant system: the impact it has on the electric grid, improved human comfort, and the efficiency afforded by thermodynamics.

Motivation 1: The Impact on the Electric Grid

Regarding the electric grid, more energy is consumed by buildings in the U.S. than by either transportation or manufacturing.

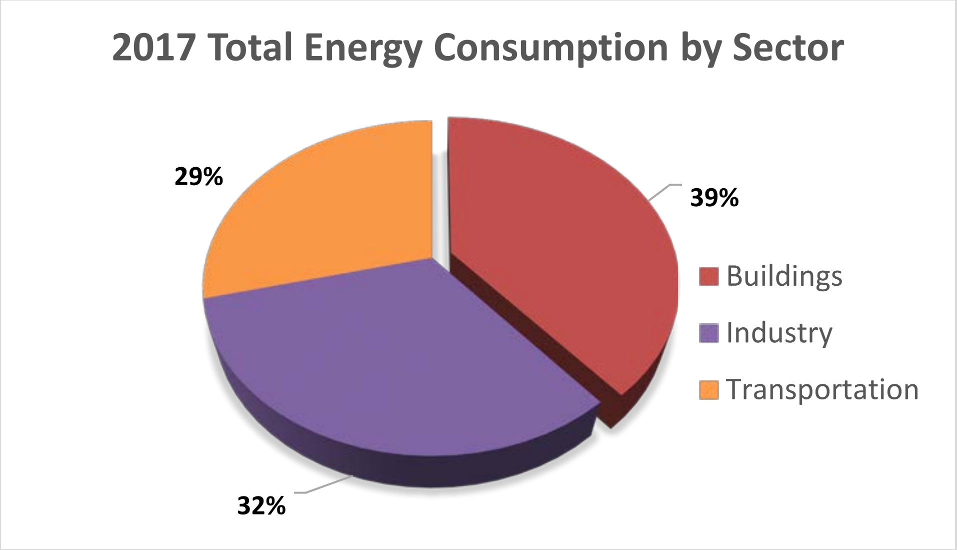

Figure 1.1: The energy use breakdown in the U.S. in 2017 [17].

While the share of energy used in buildings is large, the share of electricity is even larger. Saving building energy consumption results in fewer natural resources being devoted to energy production. The building energy consumption can be broken down further into its end uses. Figure 1.1 shows how energy in office buildings was used in 2012.

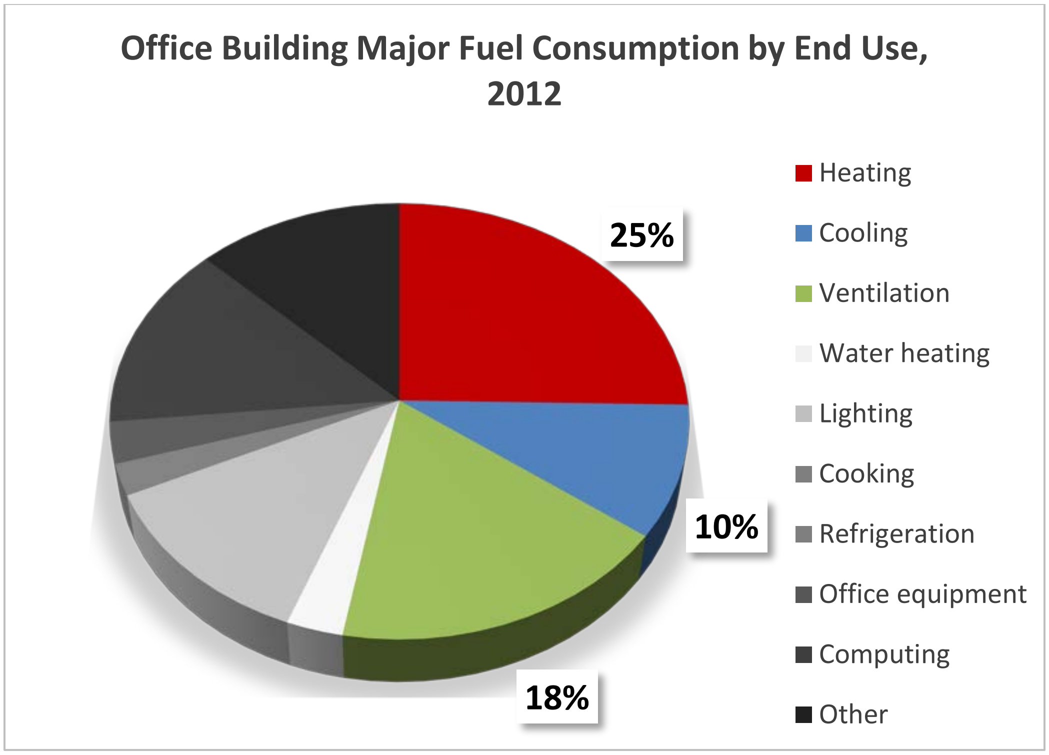

Figure 1.2: How energy is used in U.S. commercial office buildings [6].

Based on the latest data shown in Figure 1.2, over half of the energy in commercial office buildings is used on HVAC . Looking back on the entire commercial building stock, HVAC has remained the highest single category of use in buildings over the last 25 years as shown in Figure 1.3.

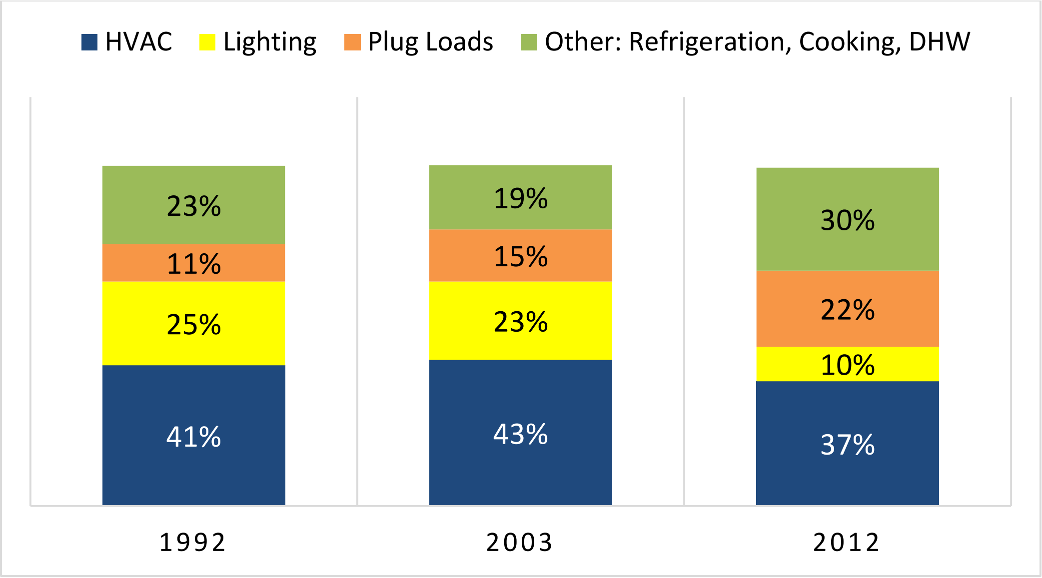

Figure 1.3: End-use of energy in buildings over the last 20 years [5].

One of the most striking details from Figure 1.3 is how the lighting consumption has changed. In 1992, lighting accounted for one fourth of all energy used in a building. In 2012, lighting accounted for only one tenth of a building’s energy and that portion continues to fall. Technological breakthroughs and lighting controls have allowed people to receive the same amount of light while using fewer watts to provide that light. Amory Lovins famously noted that people do not want kilowatts – they want services such as lighting, heating, cooling, and the ability to run their computers [18]. It is up to engineers to deliver those services with as few kilowatts as possible. The energy reduction in lighting has come about through better controls, such as photosensors and occupancy sensors that reduce the lights when they are not needed. New technology, such as LEDs, has also helped to dramatically shrink lighting energy consumption. Over the same period, the percentage of energy devoted to HVAC has remained relatively steady. Radiant systems provide a paradigm shift in heat transfer and are one way of delivering thermal comfort without requiring as much energy. Radiant systems have been shown to address better comfort metrics and reduce the typical HVAC energy in a building by 40-58% over traditional Variable-Air-Volume systems [15].

Motivation 2: The Impact on Occupant Comfort

The second motivation to use a radiant system is human comfort. While many buildings control for comfort through a thermostat setpoint alone, the governing factor in thermal comfort is actually the mean radiant temperature – the average of the surrounding surface temperatures. Pursuing efficiency and achieving environmental credentials can tempt designers and building managers to widen setpoints and push the boundaries of occupant comfort. This strategy can lead to significant savings and was even promoted by former President Jimmy Carter in his Fireside Address encouraging citizens to lower their thermostat setpoints in the winter and wear sweaters [19]. While reducing electric consumption is a noble goal, doing so at the expense of occupant comfort can be costly in other ways. There can be significant pushback against this strategy and any occupants that are too cold or too hot will not be as productive. The reason to have an HVAC system is so that it can provide comfort. Significantly widening the setpoints to reduce energy comes at the high price of lost comfort and productivity. One could consider the caves of Neanderthals to be net-zero, but few would want to regress our civilization that far.

Dr. Steve Tom compiled several studies that attempted to quantify the impact that thermal comfort has on occupant productivity [8] as shown in Figure 1.4.

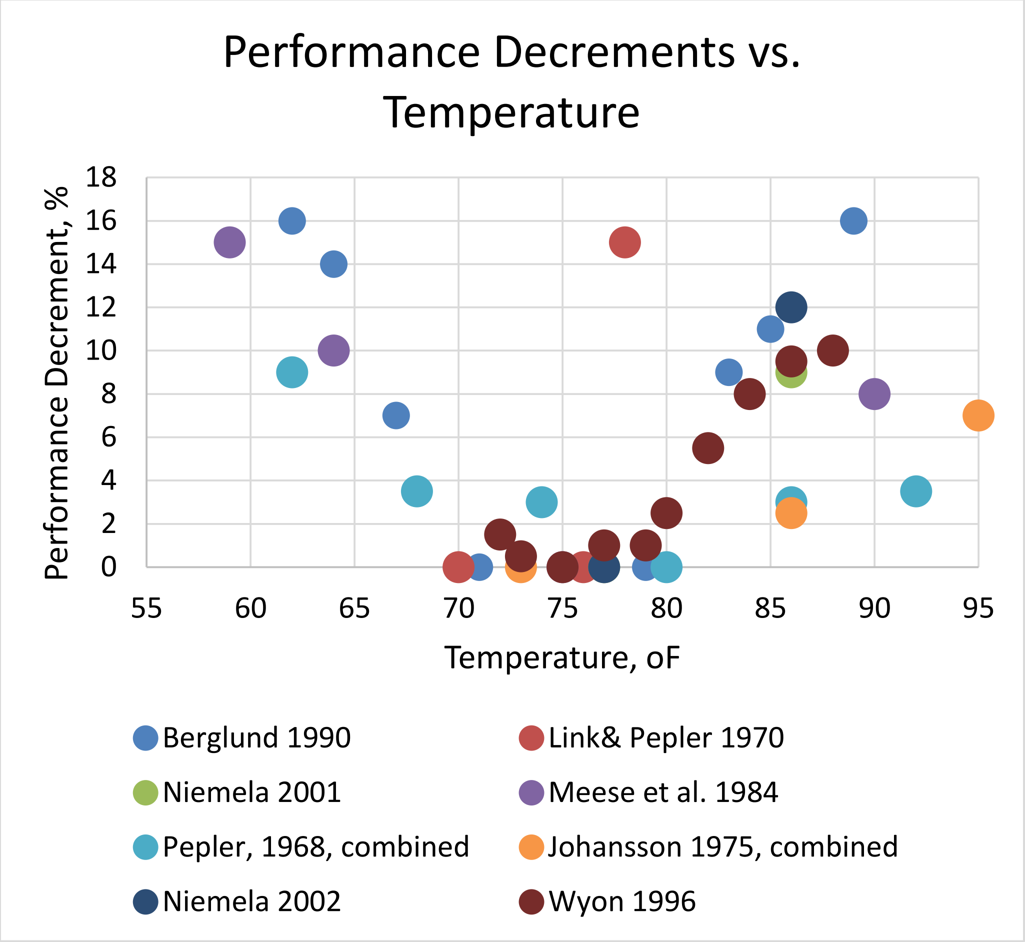

Figure 1.4: A compilation of performance versus temperature surveys compiled by Steve Tom [8].

Unsurprisingly, all the studies showed that thermal discomfort can have a devastating impact on productivity. This is meaningful because the highest expense in any building per square foot is not the construction or the utility bills, it is the occupant salaries.

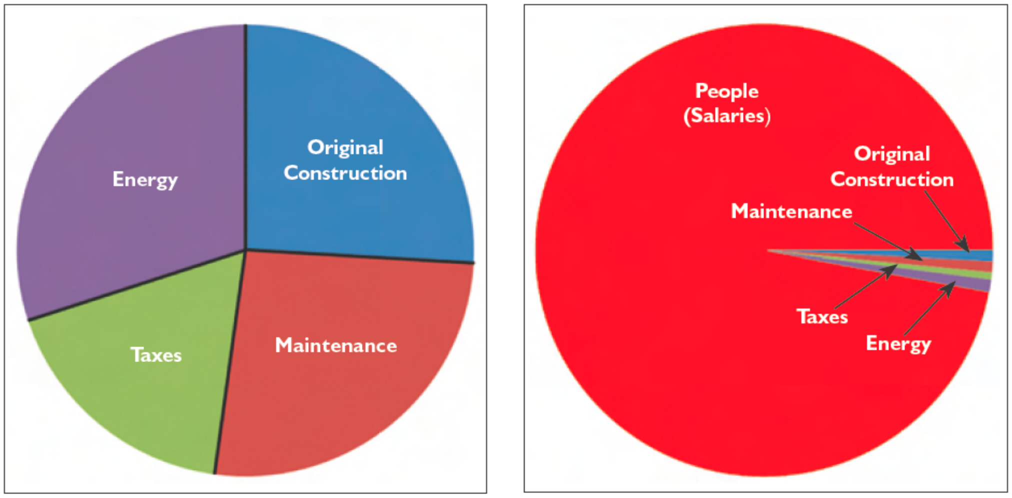

Figure 1.5: On the left is a breakdown of the life-cycle building costs and on the right is a building cost breakdown over the lifecycle of the building from Steve Tom [8].

Figure 1.5 shows how occupant salaries dwarf all other costs in a building's lifetime including construction, maintenance, and energy. From an economic perspective, the prime directive would be to maximize comfort no matter the cost of energy to maximize occupant satisfaction and productivity since salaries are the main expense. Occupant comfort is sometimes conflated with thermostat setpoints. However, this is an oversimplification that misses the nuances of how humans interact with their environmental surroundings.

There are three main ways in which humans exchange heat with their surroundings: radiantly, convectively, and metabolically as shown in Figure 1.6.

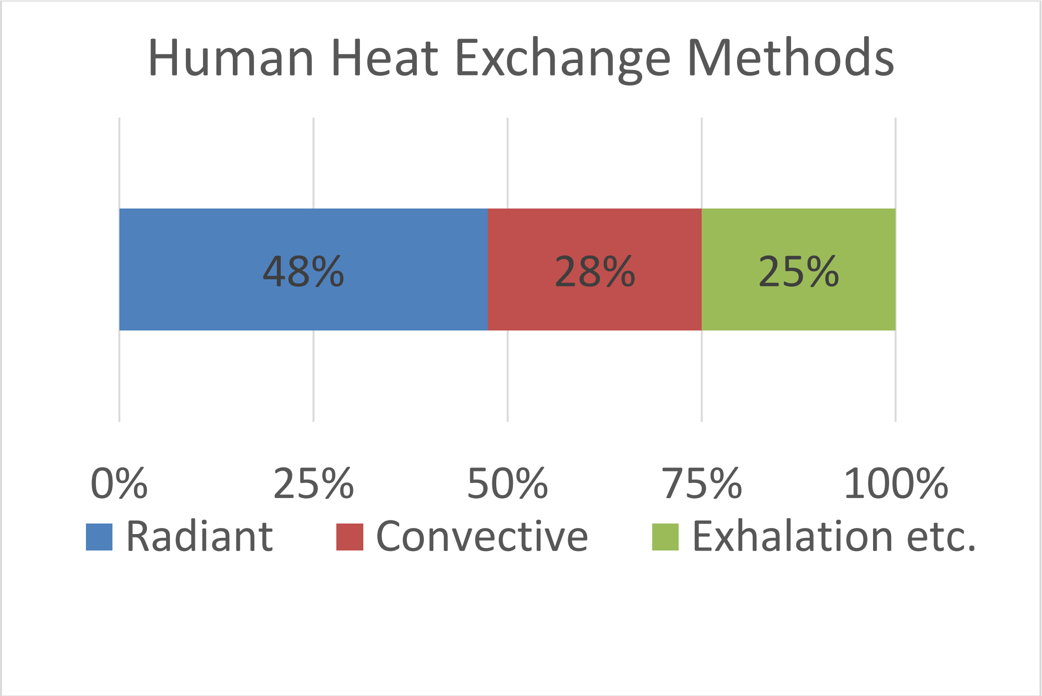

Figure 1.6: The contributing factors of heat exchange for humans from Kiel Moe [12].

We radiate heat to all the surfaces around us, exchange heat through the air that touches us, and breathe in and move in our surroundings. Of these three heat transfer methods, the most influential to our sense of comfort is radiant transfer. The surrounding surface temperatures have nearly twice the impact on our comfort than that of the air temperature. Extensive experiments by Dr. Fangar and others have helped engineers to develop general guidelines for conditions when most people feel comfortable [20]. Comfort is achieved when several conditions are all within a specified range as noted by ASHRAE Guideline 55 [21]. These criteria include the mean radiant temperature, the air temperature, velocity, humidity, clothing level, and activity. The Lawrence Berkeley National Lab has developed an online comfort tool that can provide insight into just how much the mean radiant temperature can affect occupant comfort. One example of this tool is provided in Figure 1.7.

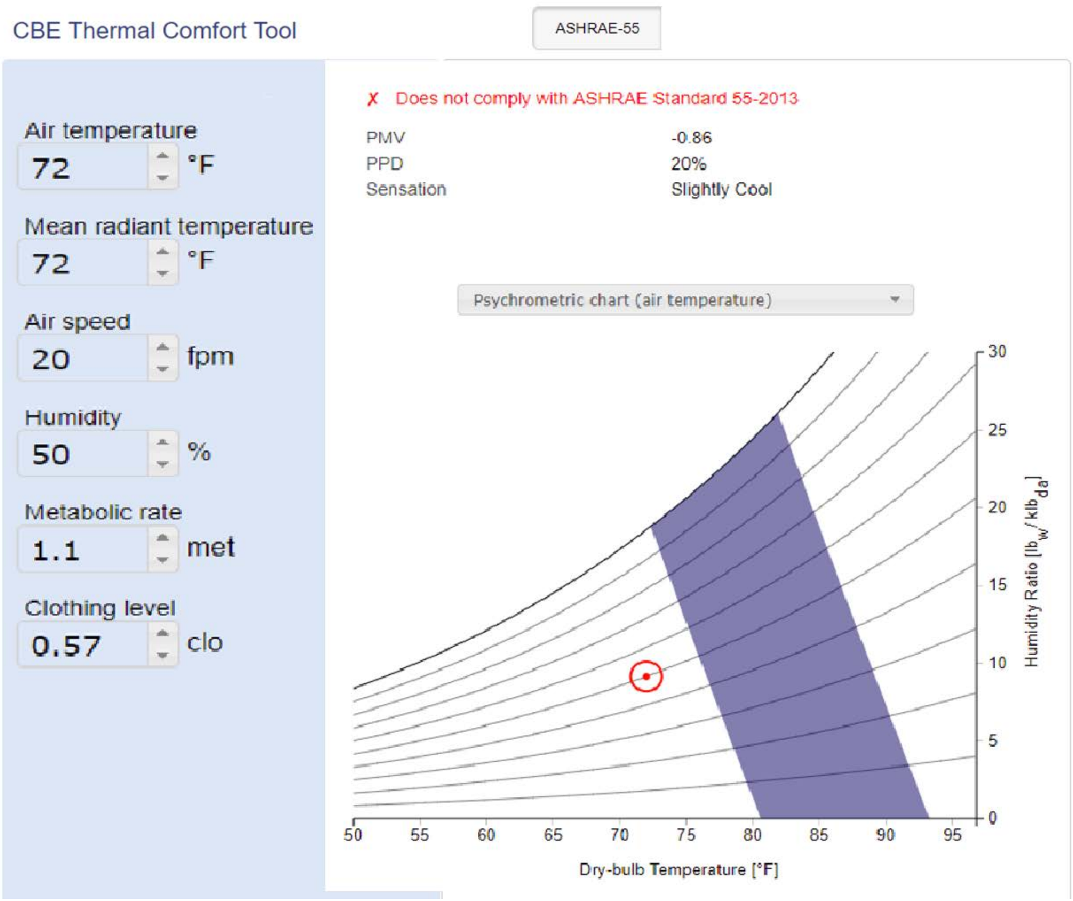

CBE Thermal Comfort Tool

Air temperature 72 °F

Mean radiant temperature 72 °F

Air speed 20 fpm

Humidity 50 %

Metabolic rate 1.1 met

Clothing level 0.57 clo

✗ Does not comply with ASHRAE Standard 55-2013

PMV -0.86

PPD 20%

Sensation Slightly Cool

Figure 1.7: An output of the CBE Thermal Comfort Tool demonstrating lack of compliance with comfort standards.

In Figure 1.7, the comfort band is represented by the purple shaded area on the psychrometric chart, while the red target represents current conditions. In this example, the air temperature and surface temperatures were set to 72^{o}F, the air speed was set at its default of 20 feet per minute, the metabolic rate was assumed to be that of someone sitting and typing, and the clothing level was selected for a person wearing socks, shoes, pants, and a long-sleeve button down shirt, i.e. standard office wear. As one can see from the chart, these conditions would lead to most people feeling slightly too cool, with about 20% of occupants expressing dissatisfaction. This provides a much clearer indication of comfort than merely relying on a thermostat’s temperature.

The operative temperature of a space is one metric that can provide a better understanding of comfort. The operative temperature takes into account both the mean radiant temperature, the air velocity, and the air temperature and provides a better criteria for conditioning a space to a person's thermal preferences. Radiant systems address the primary metric of thermal comfort – the mean radiant temperature.

Motivation 3: Advantages of Thermodynamics

The last and most logical reason to install a radiant system is the thermodynamic advantages gained by using water to transport heat instead of air. After a cursory look at the thermodynamic properties of air and water, one begins to wonder why we ever began conditioning buildings with all-air systems. Air is porous, it insulates against heat transfer and has a low heat capacity.

Water on the other hand, is extraordinarily conductive. It has a larger heat capacity and is far denser than air, allowing it to transport heat much more efficiently than air. Consider the thermal energy that can be carried in one block of each substance as shown in Figure 1.8.

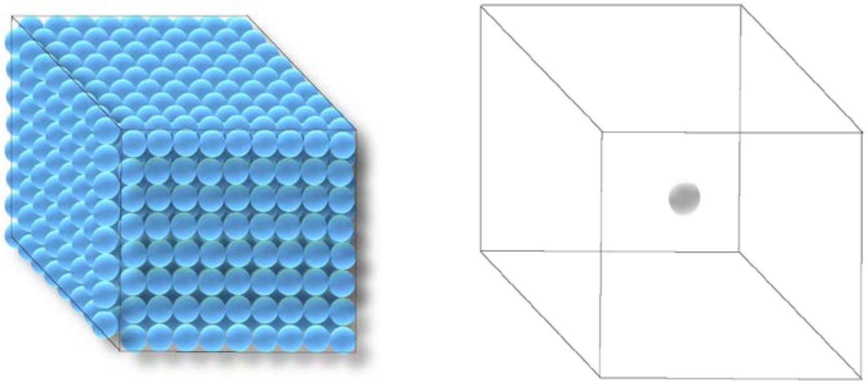

$\rho_{water} = 1{,}000\ kg/m^3$ $\rho_{air} = 1.2\ kg/m^3$

$C_{p_{water}} = 4.2\ J/g\ ^oC$ $C_{p_{air}} = 1.0\ J/g\ ^oC$

Figure 1.8: Visual representation of the density of water versus the density of air – inspired by Kiel Moe.

Based on the physical properties shown in Figure 1.8, one can see that water is 833 times more dense than air at Standard Temperature and Pressure (STP) and has a heat capacity that is four times greater than air. For every cubic foot of air, the same volume of water could transport 3,500 times the amount of thermal energy. For buildings, this means that narrow pipes of water with small pumps can add or remove much more heat than large air ducts that require huge fans.

A simple thought experiment can help illustrate how radiant systems have such an advantage over all-air systems because of simple thermodynamics. Let us say there is a room that has a cooling load of 1 ton, and the engineer wishes to remove that heat at a 10^oF ΔT as shown in Figure 1.9.

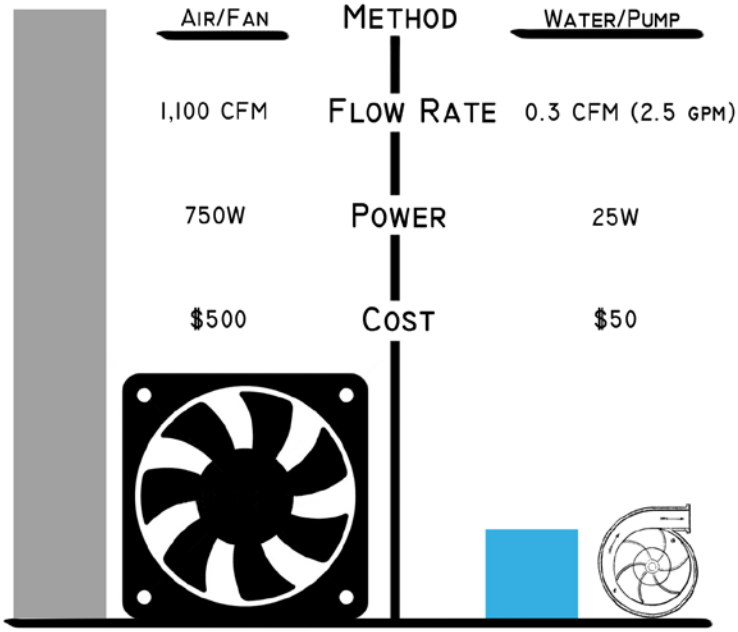

Figure 1.9: Simplified case study- provide one ton cooling based on a change in temperature at 10 ºF.

One could move this energy with a small flow of water at 2.5 gpm, or a large flow of air at 1,100 CFM. The water pump required costs about $50 off the shelf and takes 25 W to run. The air fan required costs about $500 and takes 750 W to run.

Design Implications

One of the distinct design advantages of using radiant systems is the amount of interior space that is saved by having smaller air ducts. A 1-inch diameter pipe can carry more thermal energy than a 4-foot by 4-foot air duct. Since all that is required of a DOAS in a radiant building is ventilation air, the ductwork is significantly down-sized. The smaller ductwork leads to higher floor-to-floor heights, which maximizes space and minimizes construction costs [12]. It seems that this is why the State of Washington is pushing engineers to design in this way with the new energy code requiring DOAS in office, retail, education, libraries, and fire stations [22]. With these inherent advantages, it is little wonder that at least half of net-zero high performance buildings with active HVAC use radiant systems [23]. The Lawrence Berkeley National Lab has actually created a map profiling radiant buildings [24]. These are buildings that are using radiant systems to meet at least half of their thermal loads in conjunction with dedicated outdoor air systems for ventilation.

Figure 1.10: CBE map of buildings using radiant systems.

Equipment

Unlike large air terminals and registers, the demand side equipment of radiant systems often blend in seamlessly with the architecture of a building. Radiant floors are one way of heating and cooling a space inconspicuously. A common type of radiant floor is a radiant slab. Radiant slabs are concrete floors with water pipes embedded directly in the concrete. The pipes can be metal or plastic. While corrosion can be a concern for metal pipes, the advancements in materials science has led to the widespread adoption of cross-linked Polyethylene (PEX) pipes which are often guaranteed for 50 years or more.

These PEX pipes between 1/2" and 3/4" rest on top of two or more inches of Expanded Polystyrene (EPS) foam as shown in Figure 1.11.

Figure 1.11: An image of the Viega radiant system before concrete is poured at the CSHQA site [25].

The topping concrete slab is poured directly over the PEX piping and insulating foam [25]. In addition to the floor, a building's ceiling can also be used for radiantly conditioning a space. For example, chilled beams take advantage of structural elements and can be incorporated into the air distribution system while chilled panels act as functional ceiling tiles that lower the surface temperature above the occupants. Even combining air distribution with radiant pipes can be both subtle and effective. Using a radiant system for cooling offers unique advantages in downsizing traditional equipment. Chillers can be minimized or done away with altogether and replaced by heat pumps, cooling towers, or a combination of the two. The heating equipment can be similarly downsized. While radiant systems hold a number of advantages, they suffer from some common misconceptions in the field: that they respond too slowly to gains, are too expensive, cause unwanted condensation, may not save enough energy to be cost effective, and may not have the capacity to fully condition the space. Each of these can be a valid concern and will be addressed beginning with the system capacity.

Radiant System Capacities

The first consideration for any space is how much of the heating and cooling loads of the space can be handled by the radiant system. The floor or ceiling area required for a radiant system varies by location and application. Designers should begin by limiting all internal loads first, this includes mitigating solar gains, insulating walls, adding low-wattage lights and minimizing plug loads, all things that contribute to radiant heat gain [26]. If the system capacity is not able to heat or cool the area on its own, then one can begin to increase the share of heating and cooling done by the DOAS.

Target Capacities

One unique feature of radiant systems is that their target capacities vary based on where they are installed. Radiant systems will have a higher heating capacity if installed in the floor. Any cooling system will be especially effective in an area that receives direct sunlight. However, direct solar gains can quickly overwhelm the cooling system.

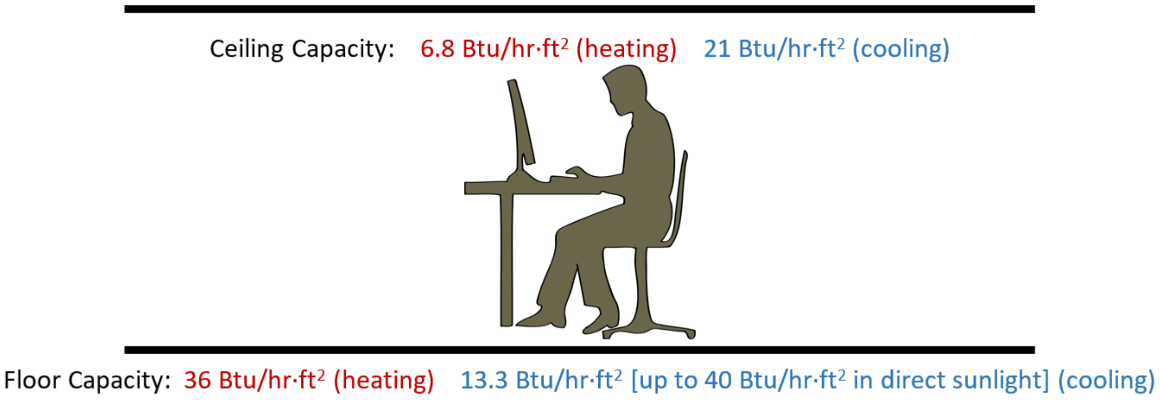

Figure 1.12: Heating and cooling capacities of radiant surfaces [27] [28].

Calculating Loads

Calculating radiant loads is extraordinarily complex. The calculations involve view factors, reflective and absorption properties for all surfaces in a room. Estimating radiant loads by hand takes hours for just one point in time. Therefore, a primary recommendation from the Radiant handbook is that designers should rely on simulation programs to calculate radiant loads [26]. Most conventional methods will fall short of the computational rigor required. While target capacities for radiant systems can be found on Figure 1.12, the best way to know whether or not the system will be able to maintain occupant comfort is to develop an energy model that is as detailed and accurate as possible so that the system is not over or under-sized. For example, the peak cooling day for a radiant system may occur in September or October when the sun is lower in the sky and direct sunlight infiltrates the floor area. This dynamic could be missed when relying on a simple air-side cooling load calculation.

Radiant heat loads and convective heat loads are like two different thermal languages. Currently most software packages are designed around convective heat loads. Among energy simulation software packages, there are two methods of calculating radiant heat transfer. The first is known as the radiant time series method. This method in effect delays the conversion of radiant gains to convective gains depending on a material's mass. The second calculation method is known as the heat balance method which takes a first-principles approach to calculating radiant heat transfer and is much more accurate at simulating radiant gains. The ASHRAE handbook of fundamentals recommends using software with the heat balance method for the most accurate simulations of radiant systems [29]. Softwares such as DOE 2, eQuest, and Trane Trace (pre-2017) use the radiant time series method, while software such as IES (VE), TRNSYS, EnergyPlus, and ESP-r use the more rigorous heat-balance method.

Response to Thermal Loading

Convective heat loads and radiant loads can operate at very different time scales [27]. In understanding this difference, the following thought experiment may provide some clarity: imagine it is a warm summer day, and a car has been parked in the sun all afternoon. When the driver gets into the car, all of the surfaces are hot to the touch and the driver may begin sweating immediately. When turning on the air-conditioning in the car, the cold air blowing on the driver's face may feel nice, but is insufficient to cool the driver. The first few minutes of the drive will be warm and uncomfortable because it takes time for the cold air in the car to exchange heat with the interior surfaces, lowering them enough for the occupants to achieve thermal comfort.

Radiant systems experience solar and internal loads immediately – as they happen. All-air systems on the other hand may experience hours of delay before these gains are registered on a thermostat. Take for example an office employee sitting in an office next to an unshaded window. As the sunlight comes through the glass, it splashes across the floor and the desk. The employee feels the warmth from the floor and desk immediately, even though the thermostat remains unchanged. In fact, it may take hours for the surfaces in the office to warm up enough to begin a convection process which will then change the air temperature. This unique load profile as shown in Figure 1.13 was captured well in a publication by Bauman and Feng which shows that radiant systems experience a much higher peak cooling load earlier in the day than an all-air system [27].

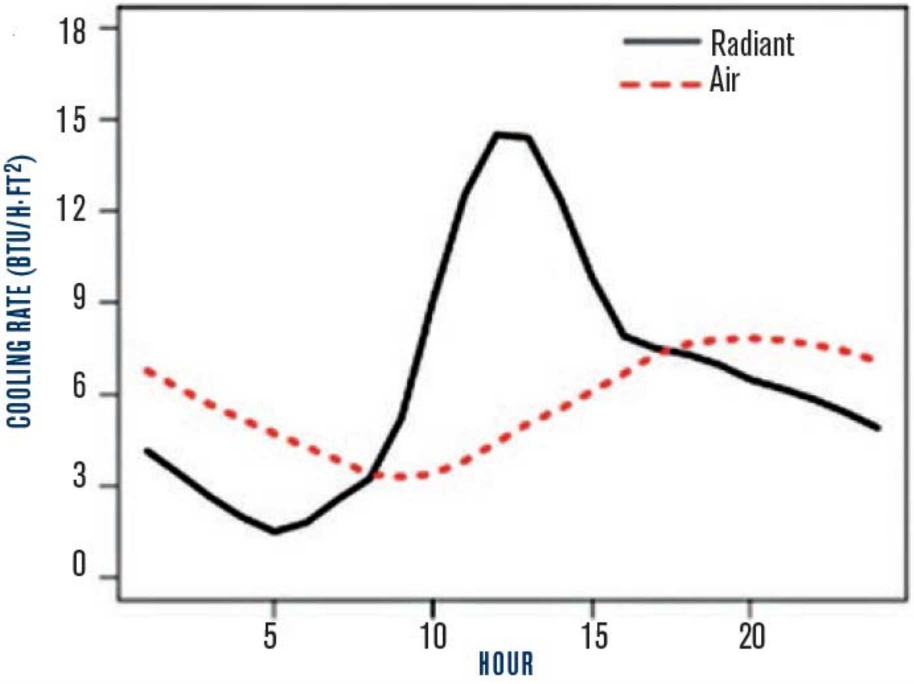

Figure 1.13: Contrasting cooling load profiles of a radiant vs an air system as developed by Bauman and Feng [9].

The radiant system must respond to internal gains much sooner than an air system. The goal of a radiant system is to keep surfaces at a reasonable temperature all the time and prevent them from warming up in the first place. If the radiant system is controlled only by the thermostat, it will respond far too late and the employee will feel too hot. Instead the radiant system must operate according to a different control strategy.

Radiant systems must respond to a higher peak cooling load than convective systems. However, just because the radiant system reaches a higher cooling peak load during the day than an all-air system does not mean that it uses more energy. On the contrary, radiant systems are far more effective at removing this heat than an all-air system. Instead, it is like having a slightly larger nail to address, but having the resources of a thermal sledgehammer.

Managing the Surface Temperatures

The goal of radiant systems is to influence the mean radiant temperature of the space. Adjusting the surface temperatures improves the occupants' operative temperature without drastically changing the air temperature. It is often surprising to many clients and engineers just how little heating or cooling the active surface requires. Commonly acknowledged surface temperature limits for heating and cooling can be found in Figure 1.14. These limits are provided courtesy of healthyheating.com, a brainchild of Geoff McDonnel, P.E. who has contributed several articles on the subject to the ASHRAE Journal [28]. These temperatures are far closer to ambient conditions than an all-air system, which must heat or cool the air to much greater extremes in order to accommodate comfort as shown in Figure 1.14. These surface temperatures are kept by flowing water through the pipes embedded in the radiant surface with a minimal difference between the supply and return temperatures. This is possible because water is so dense, and has such a high heat capacity, that a large amount of thermal energy can be transported even at a low flow and low temperature difference.

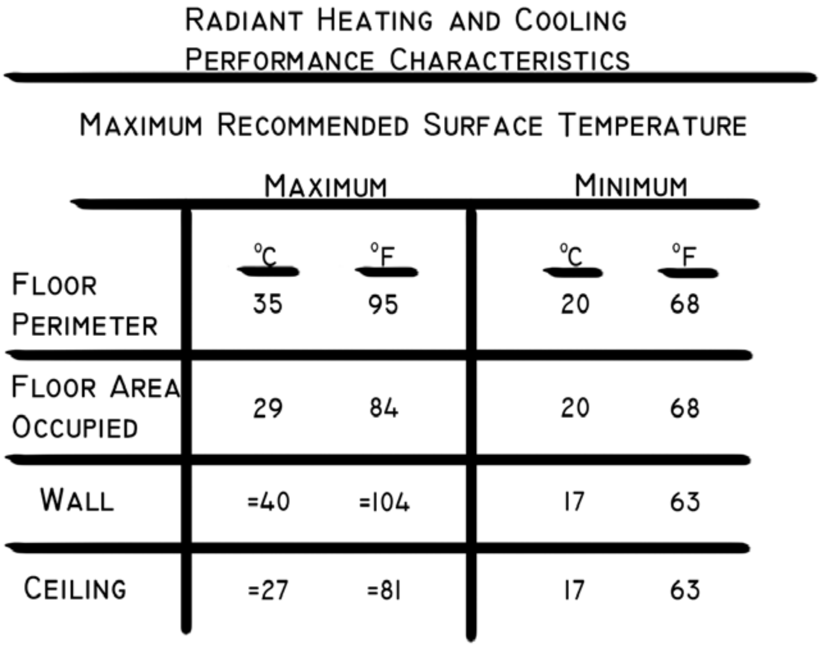

RADIANT HEATING AND COOLING PERFORMANCE CHARACTERISTICS

MAXIMUM RECOMMENDED SURFACE TEMPERATURE

	MAXIMUM		MINIMUM	
	°C	°F	°C	°F
FLOOR PERIMETER	35	95	20	68
FLOOR AREA OCCUPIED	29	84	20	68
WALL	=40	=104	17	63
CEILING	=27	=81	17	63

Figure 1.14: Maximum and minimum surface temperature limits [28] [30].

Condensation

One concern that is often mentioned with radiant cooling systems is that of condensation. It is a legitimate concern, but condensation can be easily avoided through proper controls. Studies have shown that radiant cooling can be used in any U.S. climate without risk of condensation [31]. The first line of defense is to ensure that the surface temperature never reaches the dew point by using controls that shut the system down if it ever approaches the dewpoint within 2 ^{o}F. This should rarely happen. Consider an example as shown where the indoor conditions are 78 ^{o}F, at a 40% relative humidity, at these conditions the dewpoint is 50 ^{o}F as shown in Figure 1.15.

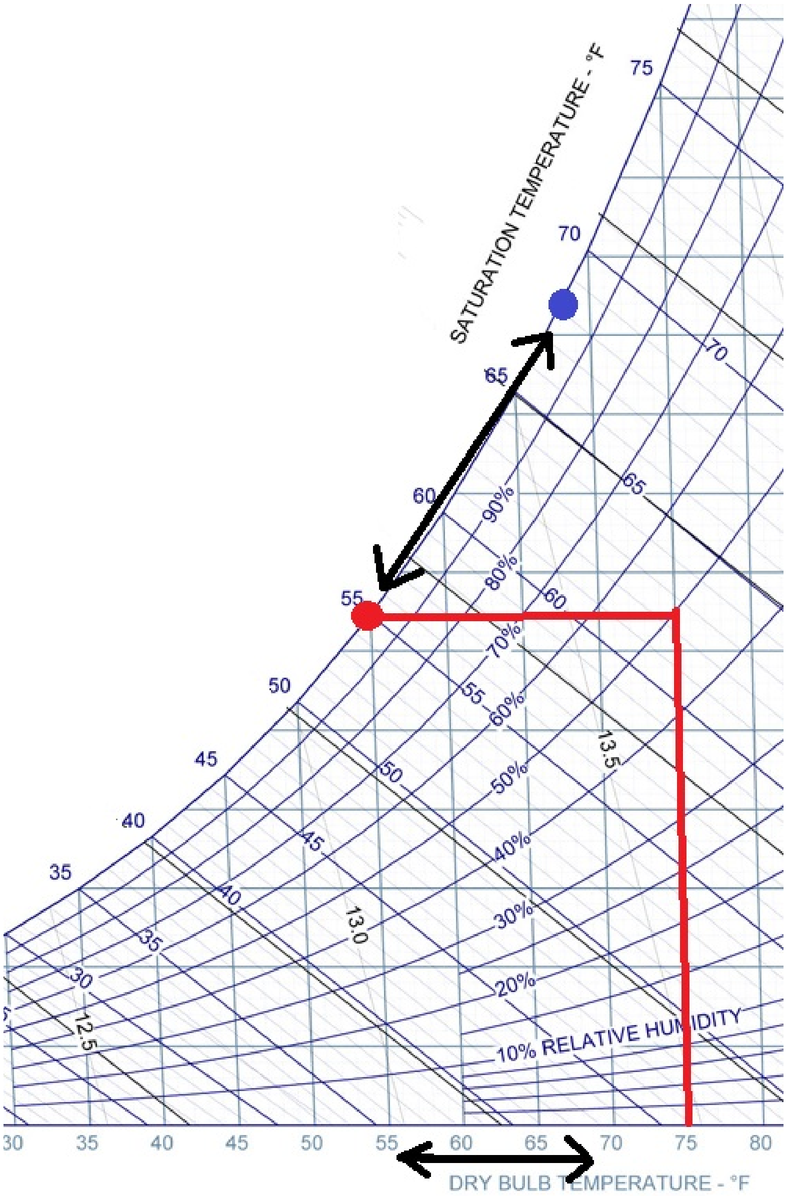

Figure 1.15: A psychrometric indicating the dewpoint in a typical indoor climate.

The dewpoint in this scenario is well below the lowest recommended surface temperature (68^{o}F). When looking at a psychrometric chart in a radiant slab situation, the dewpoint in most indoor situations is well below the lowest recommended cooling temperature (see Figure 1.15). The only time this would occur would be if there is an incredibly high indoor relative humidity. The ventilation system can limit humidity buildup. Strict limits on the maximum relative humidity can be maintained through cooling coils or dessicant wheels to remove moisture from the ventilation air for the building. Maximizing the dehumidification capacity of the ventilation system is one of the best ways to increase the effectiveness of the radiant cooling system [32].

Controls

One of the primary concerns of building engineers and operators is how to best control the system. This concern is especially true of radiant slabs, which have a high thermal mass and are slow to change temperature. Each building is a specific case and requires controls that are unique to it. One of the few universal rules in radiant controls is that they should not rely on an indoor thermostat alone [33]. Instead, the controls ought to incorporate the slab temperature and perhaps the outdoor conditions. As a general rule of thumb, a radiant slab has a lag time between the cooling signal and the air temperature response of one to two hours per inch of concrete [28]. These systems should be activated seasonally and/or based on outside air temperature – rarely is there a case where it is efficient to heat and cool the same surface within the same day [34]. Stretching the cooling demand over a longer period of time can result in a smaller plant size. In other words, radiant systems are like a large tanker that is pulled by small tugboats. It comes down to the difference between power and energy. There may be a large buildup of energy in a space over the course of a day in the form of heat. This thermal energy can be removed very quickly with a lot of power, or very slowly with very little power. A powerful system would require larger pipes and an up-sized chiller, but this removes one of the best advantages radiant systems offer in terms of small pipes and small cooling equipment. The best way to limit the peak power consumption is to use small equipment that runs for 8 – 12 hours at a time as opposed to large equipment that only runs a few hours each day [35].

An operator that manages the radiant system can adjust the cooling or heating rate by varying either the water temperature or the water flow, or both. While simultaneous management of both the flow and the temperature is possible, it can be complex and is not common practice. Generally, it is better to vary the water temperature at a fixed flow rate than to vary the flow at a fixed temperature [36]. Controlling the temperature also helps to safeguard against condensation or overheating a surface. When controlling for temperature, it is best to limit the temperature differential of the supply and return flows to 3-5 K (5-9^{o}F) [26]. An emerging control method is to rely on models to predict the best control actions. This type of control is discussed in detail in Chapter 2.

Recent Radiant Case Studies

The Infosys building in Hyderabad, India offers the clearest possible contrast between an all-air system and a radiant system. The 250,000 ft^2 building has two wings that are mirrored across from each other. The HVAC system is split exactly in half – with one wing conditioned by radiant systems and the other using an advanced VAV air-system. A profile of the building is shown in Figure 1.16.

Figure 1.16: The InfoSys building in Hyderabad India [37].

Occupancy and loads were very similar in each wing of the building. This project is set in an extraordinary climate that was cooling dominated with a peak dry bulb temperature of 115 ^{o}F and a peak dew point of 77 ^{o}F. In other words, it is extremely hot and humid. This project was designed as an experiment to see which was the better system - an advanced VAV air-system, or a radiant system. The radiant system outperformed the VAV system in three key areas: energy consumption, occupant comfort, and construction cost. The radiant system was cheaper to build, kept people more comfortable, and used 28% less energy than the all-air-system [37].

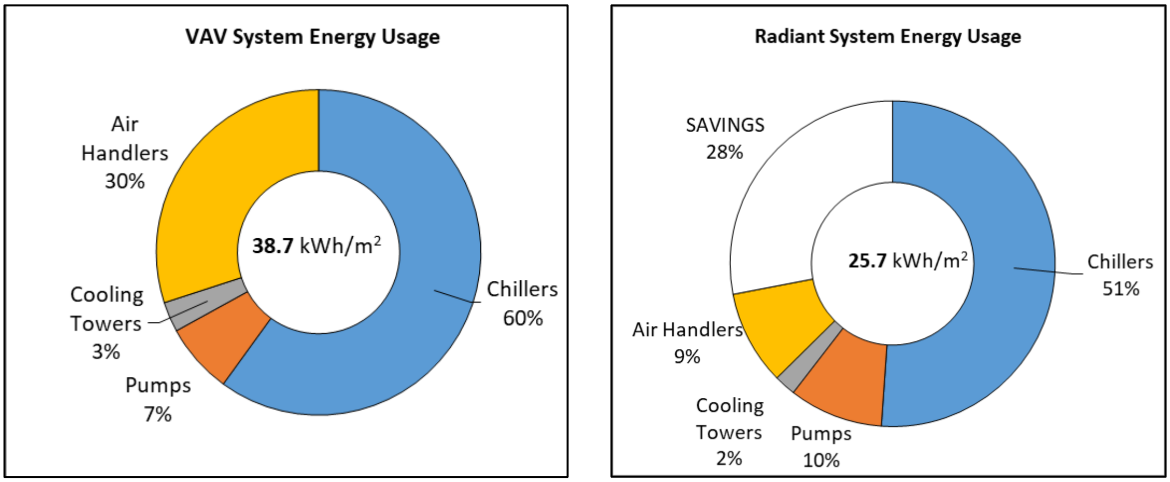

Figure 1.17: HVAC use comparison between the two wings of the InfoSys building [33]

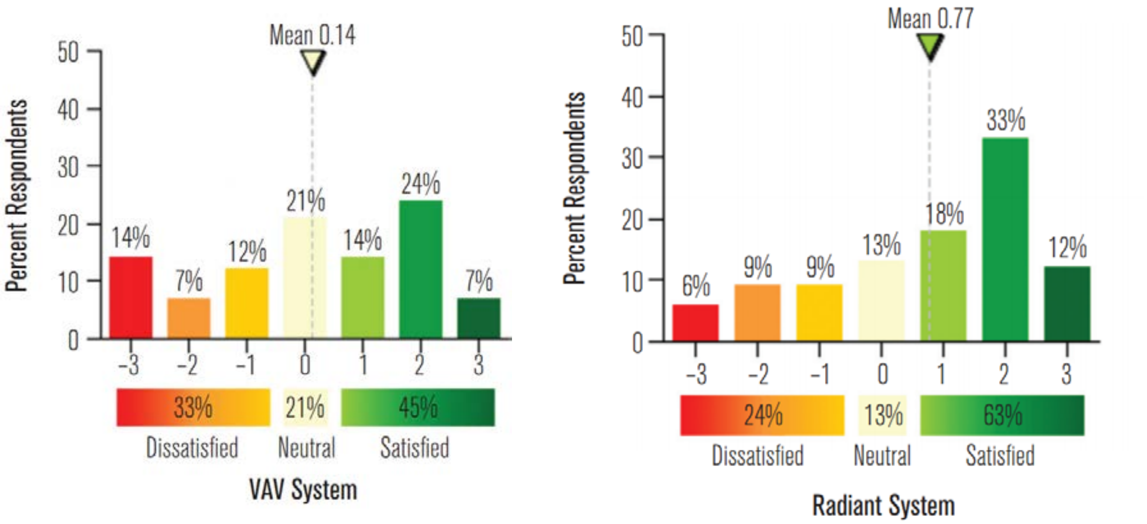

Figure 1.18: Occupant survey responses from the Infosys building conducted by CBE [33]

The Infosys building provided a real-world operational verification of both the energy savings and comfort benefits of using a radiant system.

While set in a much milder climate, the David Brower Center on the UC Berkeley campus uses advanced controls to limit its cooling energy. It provides insight into how radiant systems can be managed with minimal cooling equipment. The center’s HVAC pairs the radiant system with only a cooling tower instead of relying on a traditional chiller. The Brower Center design eliminates the need for a chiller by minimizing the building’s cooling loads. External light shelves and shading features along with internal blinds prevent high solar gains. The building also experimented with MPC to minimize the cooling tower energy consumption. Implementing model predictive control saved the facility 20% in energy expenses each year [38].

The Bullitt Center in Seattle, Washington is another radiant building that used design elements to limit thermal loads. The envelope is well -insulated and exterior blinds prevent the slab from overheating during the summer months. The building employs ground-coupled water-to-water heat pumps to heat and cool the radiant loop in the building. This six story structure has exposed ceilings so that the radiant system can positively affect the mean radiant temperature from both above and below [39].

The CSHQA headquarters in Boise, ID show that a radiant system can be applied even in a retrofit situation. The architecture and engineering firm renovated an old warehouse into a functioning design office [40]. To add a radiant system, the design team decided to lay down insulation foam on top of the existing floor, add in pex tubing, and pour the concrete over the top. The floor of the open office area serves to both heat and cool the space. The full HVAC system is described in detail in Chapter 4.

Conclusion

Radiant systems afford many advantages over all-air systems. They use significantly less energy, target human comfort better, and capitalize on the density and thermodynamic properties of water to transport heat efficiently. This results in smaller equipment but also limited capacities. These systems require careful control to respond to radiant gains or losses that are not necessarily noticed by a thermostat. In buildings where radiant systems have been installed, case studies have demonstrated increased comfort and lower utility bills. Radiant systems can benefit greatly from a more holistic control approach – particularly ones that can use predictive management methods.

CHAPTER 2: ADVANCES IN BUILDING CONTROLS

Introduction

Buildings have become increasingly complex machines that need careful control. Whereas lights ten years ago had a simple on or off switch, today many lighting systems now incorporate occupancy and photo sensors that automatically shut off or dim the lights when they are not needed. The same is true of HVAC equipment, which now may include ventilation controls based on CO_2 levels and multiple stages of cooling. The bespoke time-scales and behavior of radiant systems requires a unique control approach – particularly for radiant slabs. These systems have a long delay between initial operation and air temperature/thermostat response. This is not to say that the radiant slab is slow in cooling – it begins cooling as soon as it turns on, but it takes time for the slab to affect the air temperature.

A delay in feedback signals can result in poor control decisions. The Nobel Laureate Thomas Schelling in his seminal work *Micro-Motives and Macro Behavior* highlights the consequences of slow feedbacks and overshooting [41]. In the book, Schelling uses the example of a person standing in a shower trying to achieve the ideal temperature. In a scenario common to many, the person in the shower is impatient for the water to warm up and turns the dial to increase the amount of hot water; feeling no immediate difference, they turn the dial further until they are met with a deluge of scalding water and adjust the dial back. Schelling's example ends there, but taking it a step further, one imagines that the next time the occupant uses the shower, they run the water for a while in advance until it meets their ideal temperature before stepping in; afterwards, they make only cautious, small adjustments to the hot water dial. The same scenario plays out with radiant systems all too often because of the delay between the signal and the response when relying on a simple wall thermostat. In order to solve this issue, either the feedback loop needs to be shortened (relying on different temperature signals), or the system needs to be run in advance – avoiding major adjustments. The following chapter will outline a method for shortening the feedback loop by using better signals and running the system in advance (i.e. letting the water reach the right temperature before stepping into the shower). Model predictive controls provide one solution to this control complexity.

Model Predictive Control

Model Predictive Control (MPC) is "an optimization based strategy in which an explicit model is used to predict the behavior of the controlled plant over a receding window into the future" [42]. It is a form of hard control linked to some sort of model inside a program. One analogy is to consider a driver in a car approaching a curve. Seeing the curve up ahead, the driver does not immediately turn the wheel, but instead makes a series of small adjustments. As the car travels further, the driver has new information and can make another minor adjustment in a continuous feedback loop. In this analogy, the car is the system and the driver is the controller. MPC gained popularity in the chemical engineering industry since it is well-suited to anticipate future heating or cooling needs of a chemical process while avoiding drastic temperature swings by looking ahead and making small ongoing adjustments. This same technique can be applied to building controls.

Traditional MPC approaches rely on a mathematical derivation of the thermal loading of the building based on various parameters. This is expressed as a differential equation, sometimes called a reduced-order-model (ROM). A limiting variable is chosen for optimization, and then dozens to thousands of iterative simulations are used to determine the best approach. At the next time-step, given new inputs, the simulations are run again for the next control step. Each MPC requires a history and a forecast horizon, a way to incorporate future predictions into the next step. Because so many iterative calculations are required, the simpler the model, the faster the MPC can respond. If a new control input is required every minute, but it takes five minutes to run the simulations, then the controller will never be able to keep up and the system will break down. Therefore, complex models are rarely used in MPC schemes. Since buildings are so complex, the MPC approach has not been commonly used for system operation as a whole and has been focused instead on individual building components, such as chillers or cooling towers.

Developing predictive controls for HVAC components is an area of ongoing research, but has already proven capable of providing significant savings over conventional energy management systems. An MPC was applied to a district chiller on the campus of UT Austin and was able to save 9% of their energy usage and cost savings of up to 17% through optimizing chilled water storage to reduce consumption during peak hours [43]. Other MPC systems have shown even more dramatic energy savings of 30-70% [44].

While MPC is advantageous for single HVAC components, further advantages can be realized by linking multiple systems together. For example, control of dynamic shades affects the heating and cooling loads of a building. Combining the control of a radiant slab and a dynamic façade can lead to substantial energy savings. Radiant slabs that serve to alternately heat and cool a space operate on much longer time-scales than many other HVAC components and are heavily influenced by solar gains through unobstructed windows. This makes radiant slabs and active façades prime candidates for comprehensive governance by an MPC.

Increasing computational speeds have opened the door to new ways of using MPC. While differential equations were traditionally used, it is now conceivable to use a full energy model to run parametric simulations. A detailed energy model might take a few minutes or a few hours to run an annual simulation. When the time horizon is lowered to just a few hours instead of a full year, the simulation requires only seconds. Model Predictive Control assumes the full integration of the model and the control system. An alternative approach is to have the model run in parallel with the system instead. The author proposes the term “Model Guided Control” to distinguish this de-coupling.

Model Guided Control relies on the model only for reference and information. It can serve as a test-bed for alternative control schemes without harming the equipment. It can also be linked to internet resources such as weather forecasts without fear of opening up the actual BACnet controls to malware or other intrusions. De-coupling the model from the controls enhances safety but requires physical intervention to actually enact the control recommendations. When using a full energy model, the simulation outputs must be processed and analyzed to choose the best scenario. The research team helped develop a version of this parametric simulation and processing approach. The team was a collaboration of web developers, economists, and energy modelers. The resulting Spark tool runs a set of iterative energy simulations on a cloud server for rapid decision making, similar to an MPC. This parametric simulation Spark tool formed the pattern of the research approach to controlling a radiant slab system.

Spark: A Parametric Simulation Approach

The Spark tool the research team built is a means of estimating energy savings of office building retrofits [45]. When a building owner decides to retrofit the building, there are a dizzying number of decisions to make about which items should be upgraded and which should be left alone. It can be difficult for owners to know the best course of action for their specific building, just as it can be difficult to know the best control step for a radiant slab. The Spark Tool was developed to help with these decisions.

The Spark tool enables rapid simulation of measures on demand so that owners can know the best energy efficiency measures to implement for the next 30 years of the building's life. Work on this project formed a structure for building HVAC iterative MPC for radiant slabs. The tool is accessed by a web interface where owners may input details regarding their specific building. Based on this information, a match is selected from a database of OpenStudio models. There are a number of measures that are preset in the program, allowing for a custom set to be applied to each new project depending on its needs. Two levels of efficiency are targeted – the first level is a minimum energy savings of 35% over current consumption, the second level of savings calculated is the maximum technical potential at the site based on the Spark tool's set of measures, with their inherent performance assumptions. The tool also assesses costs and benefits of each of these targets in financial terms familiar to the building owners. This report is centered on the development of the models and the associated energy efficiency measures. The Spark allows for rapid simulations and scenario testing without requiring owners to download new software or to devote prohibitive amounts of time or resources to the endeavor.

After entering building information into a web interface, users may select various upgrade options and test different combinations of measures. A PDF report is generated with estimates of implementation costs and financial returns from an investment analysis. The way this works is by using OpenStudio Energy Plus models and applying measure scripts that alter the model before simulation depending on each user's unique selections. The user may then choose to go back and select different options to compare the costs and gains of each scenario. This workflow is illustrated in Figure 2.1.

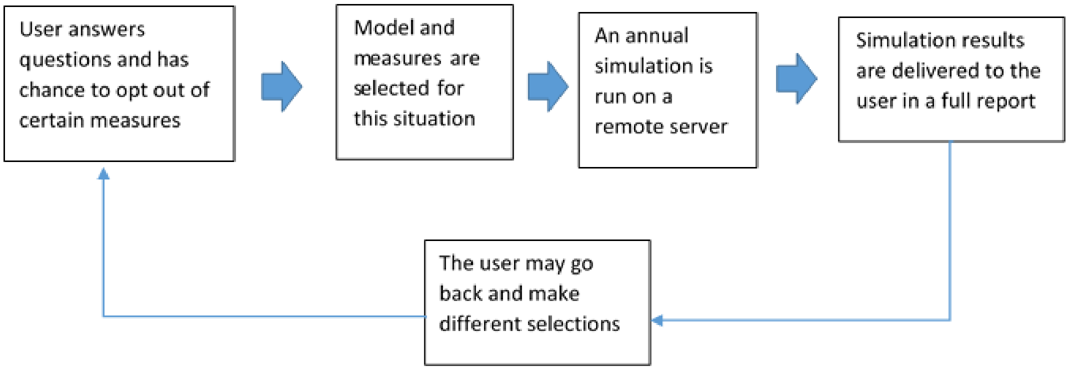

Figure 2.1: Project workflow of the Spark tool operation.

Several other simulation-based tools can also be used for determining building-appropriate retrofits and savings. The challenge of these tools is to balance modeling fidelity without making the tool complex and cumbersome for the user [46]. One method to avoid simulation run-times and complex inputs is to rely on a pre-simulated database. Such databases are used in the tools C3 Commercial, FirstFuel, and SIMIEN. Simulation-on-demand allows for a wider array of user inputs. The Commercial Building Energy Saver (CBES) tool from LBNL is a comprehensive retrofit analysis tool that uses EnergyPlus simulations on demand [47]. The CBES tool is currently designed only for small-to-medium office buildings, although it is under development to handle more building types. The Spark tool also uses simulation on-demand based on user-inputs. Unlike CBES, the Spark tool can also be used for both large and medium office buildings.

The Spark tool offers simulation-on-demand for about two dozen retrofit options and includes an economic analysis of the results. The calculations rely on three different sections: a user questionnaire, an OpenStudio model, and a set of energy efficiency measures. This is the heart of the tool that operates behind the scenes of the web interface. While it is invisible to the user, it is the basis of Spark’s energy savings and economic analysis. The models developed in OpenStudio are known as “seeds.” Each seed acts as a starting point for a different building situation. The mid-rise model is used for any building less than six stories. The team chose to use models based on specific buildings in the Pacific Northwest that were representative of the commercial building stock [48]. These buildings served as demonstration projects for this combined Energy Efficiency Measure (EEM) approach. Calibrated models were developed for these two buildings. However, because these buildings are real, they were unique and had only one HVAC system associated with each of them. Thus, based on the same loads, vintage, and geometry of the building, several different HVAC systems were modeled. Writing a measure to replace the entire HVAC system for a building can be a challenge. Therefore, different seed models were created for each type of HVAC system that would be appropriate. The schedules and lighting levels were based on a study of an office building in the Pacific Northwest [49].

The tool is intended to provide an overview of the magnitude of potential savings and is not a guarantee. Many older buildings suffer from poor operation in addition to outdated equipment. Therefore, extensive literature is offered to the users on the importance of proper building commissioning whenever considering a retrofit. The baseline model is unlikely to capture operational abnormalities, while the energy model including the retrofits is expected to have undergone commissioning and be operated according to its intended design. The tool uses the actual utility bills for the baseline energy consumption, while the simulation results with applied measures are used to estimate post-retrofit energy consumption. Based on the user's inputs, energy efficiency measures are layered on top of the baseline seed models. These measures draw upon past research [50] [51] and are specific to typical office buildings in the Pacific Northwest. The list of measures include the following:

- Increase wall insulation
- Add a secondary glazing system or install new windows
- Seal the envelope
- Reduce the Lighting Power Density (LPD)
- Install daylight sensor for perimeter zone lights
- Install occupancy sensors for lights
- Replace task lighting with LEDs
- Add occupancy sensors to plug loads
- Install optimized building controls (DDC)
- Add variable frequency drives (VFDs) to pumps on water loops
- Install a new gas condensing boiler
- Retrofit or replace the chiller
- Retrofit or replace full HVAC system
- Retrofit or replace WSHP, PVAV, or VAV
- Replace VAV with radiant heating and cooling system and a DOAS

Each measure works in a different way. For example, the insulation measure is implemented in OpenStudio by replacing the EnergyPlus script that describes the exterior wall construction to simulate the effect of adding insulation. The baseline wall is based on a pre-1980 office construction from CBECs. The CBECs baseline was used because the demonstration project had wall insulation well above code. The window measure is similarly modeled. Within each seed model are two construction types: baseline and upgrade. The OpenStudio measure replaces the baseline constructions with the upgrade definitions. Re-sealing the building envelope was modeled by replacing the EnergyPlus schedule that OpenStudio refers to for its infiltration parameter to a schedule that reduces the infiltration by 50%.

The energy performance of the building and potential savings estimated by the tool are shown in terms of the Energy Usage Index (EUI) in kBtu/ft^2. This information is also broken down into gas and electric usage, and the expense of each. The report generated by the tool can be downloaded as a PDF so the owner may study the results further without being tied to a computer. The user is encouraged to opt in to different measures and compare estimates of energy savings and financial benefits of various scenarios. The web tool lends itself to this iterative process so that minimal user effort is required.

One of the difficulties of estimating the energy savings is that some buildings may not be operating according to the design intent. This difference can lead to either over or under-estimating potential savings. For example, at one of the demonstration projects, the make-up-air unit was not running correctly, so the building was re-circulating the return air without bringing in any fresh outdoor air. This situation resulted in low energy usage for the building, but also unacceptable indoor air quality.

The user guides on the tool recommend first completing repairs and operational fixes before efficiency capital improvements are undertaken. Without correct operation, it is impossible to know how the building will respond to different efficiency measures. That is why the Spark tool is intended to estimate savings instead of offering a precise prediction. The tool does not guarantee savings, but is intended as a first step for owners considering renovations. It is meant to highlight opportunities that can then be pursued further under a more detailed analysis. Any potential project is meant to undergo thorough commissioning to ensure that the building may realize the savings from the proposed energy efficiency measures.

During beta testing, the Spark tool underwent a review by Navigant consulting. Based on this review, the tool has been revised to account for some of these ancillary loads and efficiency opportunities. The building use hours are based on an extensive study of an office in the Pacific Northwest. This tool currently has a relatively narrow scope – the renovation of office buildings in the Pacific Northwest. However, the framework can be used for a wide variety of projects by including other measures, locations, and control algorithms. The parametric simulation of energy models can serve as a structure for an MPC.

Conclusion

Model predictive controls provide a means of managing radiant systems that are complex and have slow temperature changes. Creating a model for MPC traditionally relied on complex derivations performed by a controls specialist. Now, full energy models developed by architects and engineers can be used for iterative short-horizon simulations. This workflow has many advantages as energy modeling is growing in demand. Energy modeling is required for certain rating systems, such as LEED certification, as well as certain building codes like California Title 24. Energy modeling is now recommended by ASHRAE as the preferred means of sizing equipment.

Unfortunately, after construction and/or LEED certification, many of these models are discarded. The Spark tool the authors helped to develop provides a method of running many iterations for decision making by relying on these full energy models. It compares the outputs and selects the best scenario. Bringing these approaches together allows for a highly detailed energy model to be used as the basis of MPC – bypassing the need for a controls specialist to design a bespoke ROM. Instead, the same energy model can be used for sizing HVAC equipment, obtaining LEED points, and designing advanced predictive controls. After the creation of a model, communication can be established between that model and the building controls. The following chapter details the process of connecting an energy model to a building controller.

CHAPTER 3: COMMUNICATION BETWEEN BUILDINGS AND ENERGY MODELS

Introduction

The previous chapter discussed the benefits of predictive control and the creation of a framework that can support using full building energy models to guide control decisions. The energy models can be used to make control selections based on parametric simulation outputs. The following research outlines the process of connecting a full energy model to a building's control hardware in a co-simulation mode. Using a model to contrast or correct control algorithms is referred to as virtual commissioning. Several research teams have worked to use co-simulation for virtual commissioning [52] [53] [54] [55] [56]. The following chapter builds on this research, documenting the procedural steps, challenges, and benefits of using co-simulation as part of a virtual-commissioning process.

Virtual controller commissioning links physical building control systems to energy models in a co-simulation and analyzes their performance over a full year (i.e. the simulation of a full year). If virtual commissioning can be developed, refined, and commercialized, it could become an attractive service that utilities can offer to ensure efficient use of energy intensive building systems. Using benchtop simulation tests allow for control experimentation and optimization without putting either occupant comfort or equipment at risk.

Optimal controls are essential for building efficiency. This is particularly true of radiant systems. While radiant systems ideally save energy, if the controls are out of tune, the building's air-side system may work against the heating or cooling of the radiant surfaces resulting in simultaneous heating and cooling, using even more energy than a conventional system. Radiant slabs present the greatest challenge in this regard because the feedback from the thermostat can have a significant delay and if that is not accounted for, the resulting building will be both inefficient and uncomfortable.

For every building’s Energy Management Systems (EMS), there are times when the controls either malfunction or are out of tune. Control adjustments during the first year of operation for any new building can result in a difficult and unpredictable environment for the occupants and unusual energy consumption profiles. Older buildings also suffer from control settings that have fallen out of tune. Adjustments are required after any change in occupancy or usage. Controls might have been set to a factory default which may not be appropriate for the climate where it is installed. Although the operators at the site may have changed the settings, these may no longer be adjusted for the season or the current loads. It then falls to the building operators to correct these settings.

Building operators often have multiple responsibilities and setpoint adjustments are driven by comfort complaints. Because of these demands, more subtle control malfunctions can go undetected for long periods of time. Control problems can incur exorbitant and unnecessary energy costs. Commercial buildings operated in an unintended manner can increase energy consumption by 20% compared to the intended design [57]. Therefore, a good commissioning process is essential for the proper operation of a building.

Commissioning improves both comfort and productivity in buildings [58]. While commissioning is beneficial, it is challenging to cost-effectively and continuously analyze and tune controls for a building. Every building and control system is unique (especially buildings with radiant systems), so traditional commissioning takes a long time and is expensive. Virtual commissioning through co-simulation is one technology that can assist in this control tuning effort.

Several research teams have shown that it is possible to commission a building using an energy simulation [53] [55]. This research becomes increasingly relevant as building energy modeling becomes more widespread. The connection of real building hardware to ideal software models illuminates areas of control optimization that can save energy while maintaining occupant comfort. Virtual commissioning can be used to test advanced control strategies and is an approach that can be applied to buildings at all stages of life from start-up to operation decades after construction.

A Review of Co-Simulation with Energy Models

There is a significant body of research on using energy models to commission buildings. In one study, a research team relied on a calibrated EnergyPlus model to identify abnormalities in a building's current operation in a monitoring-based commissioning [56]. While monitoring-based commissioning compares the simulation results and the building usage side-by-side, other researchers have used a method to link the building hardware directly to an energy model [54]. A team of researchers at Lawrence Berkeley National Laboratory (LBNL) have pioneered work in co-simulation through the development of a software platform that enables communication between an energy model and a physical controller [59]. This software platform: Building Controls Virtual Test Bed (BCVTB) serves as the "middle-ware" that both the energy model and controller can communicate across [60]. The way that BCVTB transfers information is succinctly described by Dr. Wen:

> "At the beginning of each time step, BCVTB blocks all the co-simulated programs and performs data exchange. Each program sends its outputs to BCVTB and gets its inputs from BCVTB through Simulator socket writing and reading respectively. As soon as the data exchange process is completed, BCVTB unblocks the operation of the programs, and each program calculates and generates new outputs during the remaining of the time step based on the newly acquired inputs. The new outputs will then get exchanged at the beginning of the next time step. The same process repeats until the specified simulation time duration is achieved [61]."

The information exchange between the model and the controller is facilitated through a communication protocol known as Building Automation and Control Network (BACnet) protocol. The controller is connected to the model using two "actors" (programs) inside BCVTB: the BACnetReader and the BACnetWriter. These actors rely on the open-source BACnet Protocol Stack [62] that downloads with BCVTB. BACnet signals comprise a series of digits that are difficult to understand without any context. The BCVTB actors BACnetReader and BACnetWriter, use associated XML files that provide syntax for the numbers so the user can more easily understand how to send or request information from the controller [60]. In this way, simulated variables such as temperature and fan status are provided to the controller in place of signals that would typically come from physical thermostats or other analog inputs.

Controllers communicate using BACnet through one of three different connections: Ethernet, Internet Protocol (IP), or Master Slave Token Passing (MSTP). BACnet Ethernet and BACnet IP can both communicate through an Ethernet cable from the controller to the computer, while MSTP communication relies on a shielded 22-gauge two-wire cable. Alternatively, BCVTB can be used to link energy models with analog signals by using a particular analog to digital converter. Using purely analog signals is an alternate pathway of model-controller communication. Previous work on model/analog data transfer has also been performed at LBNL [63].

Several researchers have used BCVTB in proof-of-concept demonstrations [64] [56], but these have typically focused on live co-simulation. In Pang's research, the energy model was built in EnergyPlus and additional sensors were added to the EMS system. BCVTB was used to connect the model to system signals and weather at the site; they took care to synchronize the BACnet signals to real-time. Later, this same research was continued at the site as a live-fault detection strategy [52]. This is similar to the approach of Li [55] who used BCVTB for fault detection diagnostics by sending all logged information to a database for further analysis. Li's research used this database collection to gather information that could be viewed in real time to assess the building's performance and identify any anomalous behavior.

In the present study, we focused on using co-simulation to rapidly commission a controller. For this reason, we used OpenStudio to quickly generate our energy model. We did not wish to pursue a live connection as had Pang [64] and Dong [52], but neither did we wish to leave the model and controller completely separate as had Wang [56]. Instead, we used a model-to-controller connection and a full year's worth of simulation data to identify controller errors. We reproduced and adapted a commissioning process closest to Li [55]. What we did differently was run remote commissioning with no live connections to compare responses between different pieces of control hardware from different manufacturers. This provided some reassurance to the facilities team, who worried that a live connection might pose a security or operational risk. We focused our research testing on a single controller at a remote location using historical weather data for inputs. We ran annual simulations on different remote controllers instead of relying on live data feeds. We pursued a strictly digital communication pathway in our research because the number of signal inputs and outputs we were considering was greater than the number of analog ports available to us on our A/D converter hardware. This remote testing of full simulations could allow control companies to pre-test their controls in a virtual environment, limiting expenses and site visits.

We developed an energy model of a three-story office building and acquired a duplicate of the building's air-handling controller. We used the model to understand the current control settings by simulating both typical design-days and full years of local historical weather from a remote location. The model and controllers of different manufacturers were used to explore the EMS logic and test new settings. From these findings, we pursued a remotely-based workflow around this technology. This virtual commissioning at an operational site enabled us to assess the commercial appeal of this process and helped us identify several barriers that might hinder widespread adoption.

Test Site and Equipment

The site selected for this project was a three-story building that serves a mix of classrooms and offices. The building's HVAC equipment consists of a typical Variable Air Volume (VAV) system with two large Air Handling Units (AHUs). One AHU serves the basement and a server room, while the main AHU provides ventilation, heating, and cooling to most of the building. The first step of the research was to develop a calibrated energy model of the existing building in the OpenStudio application for EnergyPlus. Calibration was performed by comparing the building's monthly electrical consumption for the calendar year of 2014 to the results of the EnergyPlus model simulated using an

Actual Meteorological Year (AMY) file from a nearby airport for the same timeframe. This AMY file was provided by the company Weather Analytics. The building operators provided us with hourly load and consumption information which we used to augment the calibration process. Once we calibrated the building energy model, we connected it to BCVTB. We linked the energy simulation to a controller and then studied the differences between this, and how EnergyPlus would simulate ideal control actions, to correct and optimize control settings.

The building's main AHU is responsible for much of the building's heating and cooling. Therefore, we focused on commissioning its controller. A layout of the system and main control points are displayed in Figure 3.1.

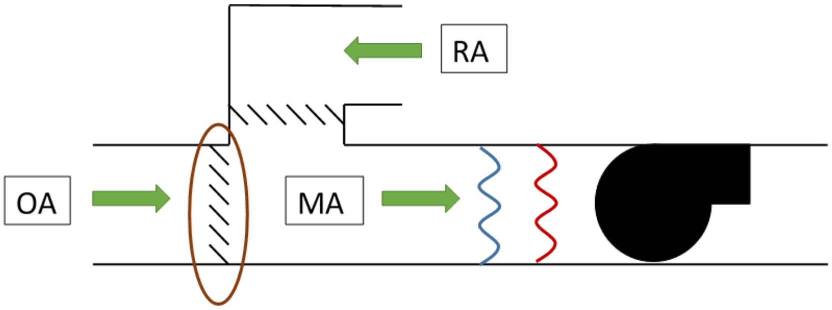

Figure 3.1: The test building's AHU layout: monitored data points are in boxes – return air (RA), outdoor air (OA), and mixed air (MA) temperatures. The control action point (the outdoor air damper) is circled.

This AHU is controlled by an Alerton VLCA-1688. We tested the virtual-commissioning process by focusing on one aspect of the AHU's operation: the control of the outdoor air into the building. We could have looked at optimizing discharge air temperature, chilled water temperature, the supply rate, or any number of other systems controlled in the building, but chose just one example for this study for the sake of simplicity. The building’s main AHU is equipped with a damper that can regulate how much outdoor air is being brought into the building. It is managed by a control device known as an economizer that can open or close the damper to bring in more outside air if it is deemed economical to do so. In theory, the economizer will bring in more cool air from outside if the outdoor temperature is lower than that of the indoor temperature so that free cooling is provided to the spaces without relying heavily on the chilled water coils. This can save the building a significant amount of energy if the control is employed correctly, but it can also cost the building more energy if the controller setpoints are not optimized. This is analogous to a radiant system which can save energy if operated correctly, but have an adverse effect if not controlled well.

Establishing Communication Pathways

When connecting the energy model to BCVTB we learned that the software is sensitive to both Java versions and path locations. Our team found it easiest to run our application using Windows 7 with 64-bit operation, and to ensure that our version of Java was also 64-bit. We also amended the Windows Path variable by appending the location of the Java and EnergyPlus installations. Dr. Nouidui provided us with assistance in navigating the installation through an online forum [65]. After adjusting the path variable and ensuring version consistency, we were able to run an EnergyPlus example file in BCVTB that downloads with the application.

We then connected our particular energy model to BCVTB. We exported our OpenStudio file to an EnergyPlus text file (idf) and opened this in a text editor. Next we added the external interface blocks of code to the end of the file. We imported and exported other variables from the EnergyPlus model to BCVTB by editing the energy model's text file and adding a short block of code for each variable as shown in the BCVTB user's manual [66]. The variables we chose for this study included the outdoor air temperature, the outdoor air damper position, mixed air temperature, and return air temperature. We selected these variables because they are the inputs required of the controller. As an initial compatibility check, we controlled the setpoints in the model through rudimentary logic within BCVTB. Once simulation results tracking a particular variable of interest showed that editing controls within BCVTB were indeed influencing the model in an expected manner, then the research moved on to establishing the controller connection with BCVTB.

Connecting the controller to BCVTB

Establishing direct model-to-controller communication was performed using a stand-alone laptop that could connect to the controller on a Local Area Network (LAN). Connection from the EnergyPlus model to BCVTB was entirely virtual (occurring within the computer). The controller on the other hand, was a separate piece of hardware that was joined to the computer through an Ethernet cable. The first controller tested was the KMC Flexstat. The Flexstat communicates through BACnet MSTP only. In order to communicate with BCVTB, the Flexstat used an MSTP connection to a MACH-ProWebcom router that translated MSTP to IP. The IP connection was then set up on the laptop through creating a LAN and changing the laptop's IP address to one that was compatible with the new network. This enabled the laptop to recognize the controller and signals could be sent back and forth to the controller. These signals were tracked using a software called Wireshark, which recorded every piece of communication between the controller and the laptop. Using this, we could see when the controller was recognized by displaying when signals were successfully sent and received between these devices. This set-up is shown in Figure 3.2.

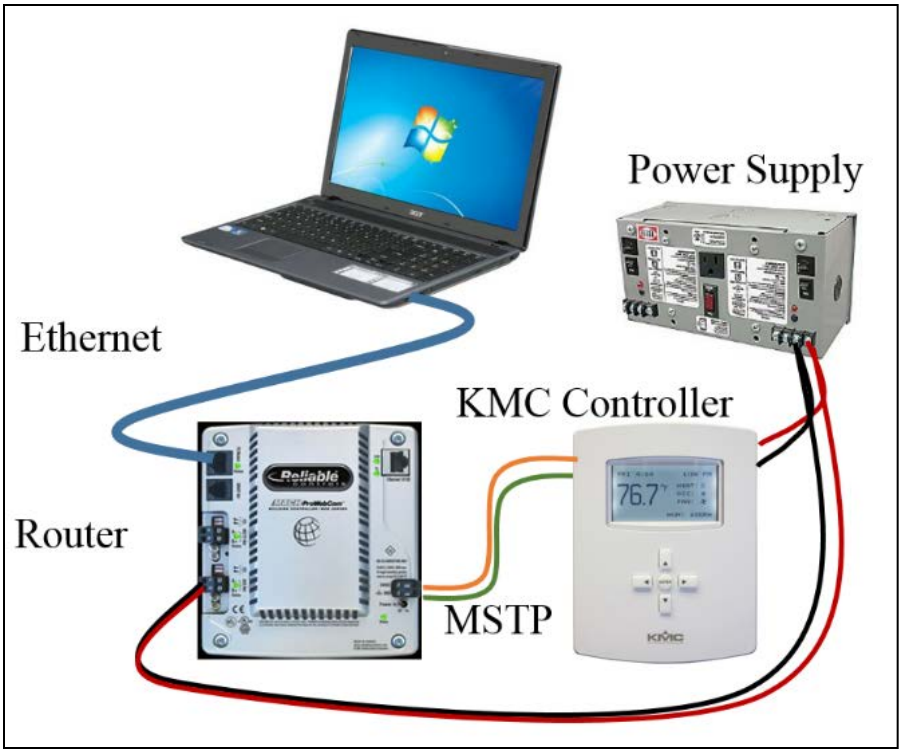

Figure 3.2: KMC controller to energy model set-up.

Both the KMC controller and the ProWebcom router required the same 10-Volt power supply. The router was the translator between the laptop and the controller, which could take MSTP signals to and from the controller and LAN through the Ethernet connection to the computer. The connection to Alerton controller (a VLCA-1688 model) as shown in Figure 3.3 proved much simpler to set up since it could communicate through Ethernet, IP, or MSTP.

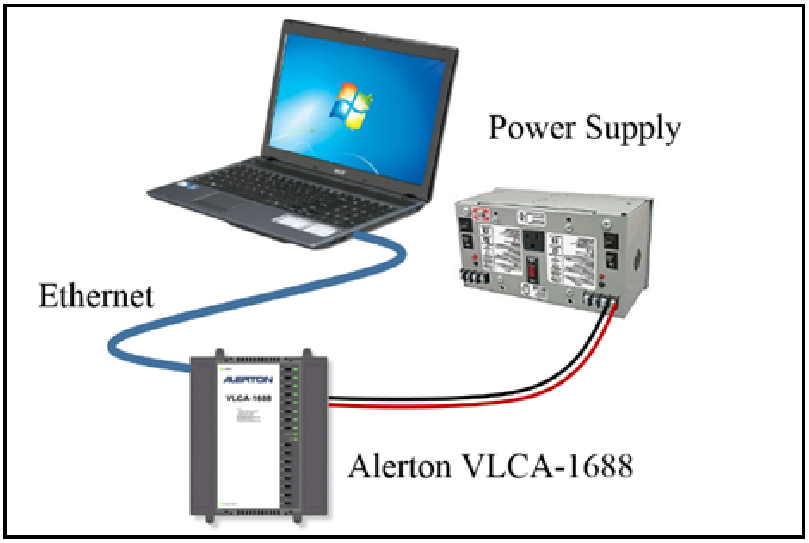

Figure 3.3: Alerton controller to energy model set-up.

We set the Alerton connection to BACnet IP so it could communicate with BCVTB directly through an Ethernet cable as shown in Figure 3.3 without a router in between the laptop and the controller.

Once we learned how to set up these devices, we studied the air-handler control settings from three different perspectives: EnergyPlus, a KMC Flexstat device, and an Alerton VLCA device identical to one at the site. We had to contend with the fact that these two controllers had very different input and output signals as shown in Table 3.1. The objects in Table 3.1 are the BACnet signals and their associated meanings. Without a full understanding of the control variables and their associated BACnet signals, it is impossible to understand and adjust the controller.

Table 3.1: Controller input and output signals

KMC -Flexstat		**Alerton VLCA-1688**	
Inputs	**Objects**	**Inputs**	**Objects**
RA Temp	AI-7	Freeze Stat Status	BI-8
OA Temp	AI-4	Fan Status	BI-5
MA Temp	AI-3	VAV's in Warmup	AV-40
Discharge Air Temp	AI-2	OA Temp	AV-0
		Econo lockout temp	AV-56
		OA Temp	AV-0
		Return Air Temp	AI-2
		Min OSA %	AV-65
		Mixed Air Temp	AI-1
		Mixed Air Max Stpt	AV-60
		Mixed Air Min Stpt	AV-61
		Cooling Signal	AV-10
Outputs		**Outputs**	
OA Damper	AO-9	Econo Command	AV-26
		Econo % Reversed	AV-29
		MA Damper Cmd	AO-4
		OA Damper Cmd	AO-3

Initial Commissioning

Once we connected the energy model to a controller, we contrasted its control sequence with the EnergyPlus default settings. These default settings in EnergyPlus functioned as an initial baseline. We recognize that these defaults are not necessarily meant to be ideal values, and these settings at any site need to be verified by a practitioner aware of the climate and codes. Yet, they served as a useful starting point to contrast the current controller settings; any major differences we encountered merited further investigation. We ran a one-month simulation with each economizer situation having the same settings: a desired supply temperature of 55^{o}F and a high-limit lockout temperature of 70^{o}F. By selecting the output variable to be the damper position, we could compare the behavior of the three different control options: EnergyPlus, the KMC Flexstat, and the Alerton VLCA model. The results revealed that all three control sequences behave similarly but at different scales as shown in Figure 3.4. The EnergyPlus default control setting opens the damper the most and the Alerton device follows the same profile, but is more conservative in its economizer usage. The KMC device is the most conservative and uses the economizer the least. When running the simulation for the month of June, the real controllers allowed in less outside air than the EnergyPlus defaults – even when they shared the same settings. When the energy model was tied to a physical controller, the simulation results indicated higher energy use because there was less use of the economizer. The lockout temperatures on each controller, and in EnergyPlus, were set to the same value, but each controller allowed in a slightly different magnitude of outside air. The stand-alone EnergyPlus model opened the damper the most, the Alerton controller opened the damper a little less, and the KMC opened the damper the least. We were encouraged to find that the three controllers behaved similarly, but through this research we realized the actual internal PID tuning of the controllers had a significant impact on how effectively they would use the economizer.

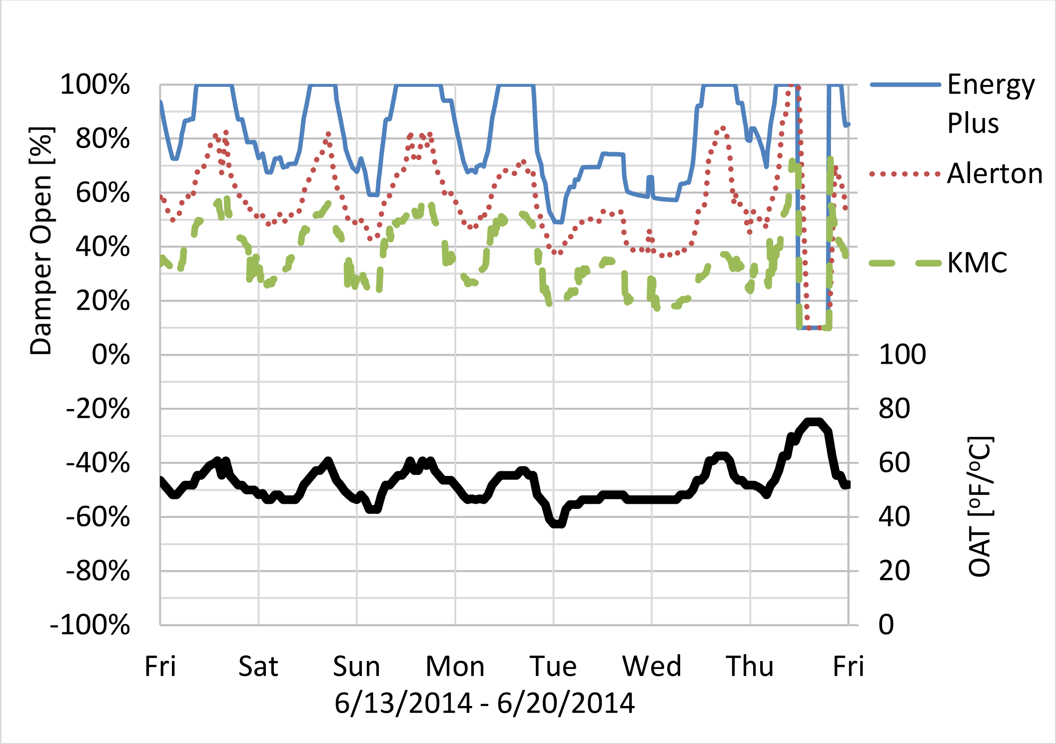

Figure 3.4: Snapshot of controller behavior during a simulated week in June.

After studying Alerton VLCA field controller settings (the actual controls at the site), we learned that the building is currently operating its main air-handler with a high-limit temperature lockout on the economizer of 60°F. Anytime the outside air climbs above 60°F, the economizer forces the HVAC system to use mostly return air from the building and only a minimal amount of outside air for ventilation. The supply air to the building is set at 55°F. The economizer ideally balances the return air and the outside air to a mix that requires the least amount of conditioning to reach 55°F. Since the economizer was set to lock out outside air above 60°F, the building loses the opportunity to use any outside air between 60°F and the return air temperature, which is typically at 76°F or above.

In order to investigate this savings opportunity, we ran three different annual simulations focused on outside air damper control. We ran each scenario using a Typical Meteorological Year (TMY) weather file for this building site. In the first case, we simulated a typical year with current controls by relying on the Alerton VLCA-1688 with the exact logic in use at the site to decide when to operate the economizer. Next, we adjusted the controller's settings to increase the maximum economizer lockout temperature to 76^{o}F and ran the simulation again. We found that it would make a difference of over 56,000 kWh (4%) per year as shown in Figure 3.5. Next, we performed further controller tuning by enabling the controller to modulate the supply air temperature based on internal zone temperatures, thus increasing the use of the economizer. The supply air modulation was disabled at the site at the time. Enabling this setting in combination with raising the economizer lockout temperature resulted in an estimated savings of 12% of the building's total energy use per year. Once again these savings were for adjustments of one control point only on a single piece of equipment. Significantly more savings could be found by further tuning and expanding the controls research beyond the economizer to other aspects of the building control, such as discharge temperature, plant loop temperatures, and setback temperatures. Implementing the simulation-based commissioning on all of the highest priority control points within the building could easily reach the 20% annual savings estimate put forth by Westphalen and Koszalinski [57].

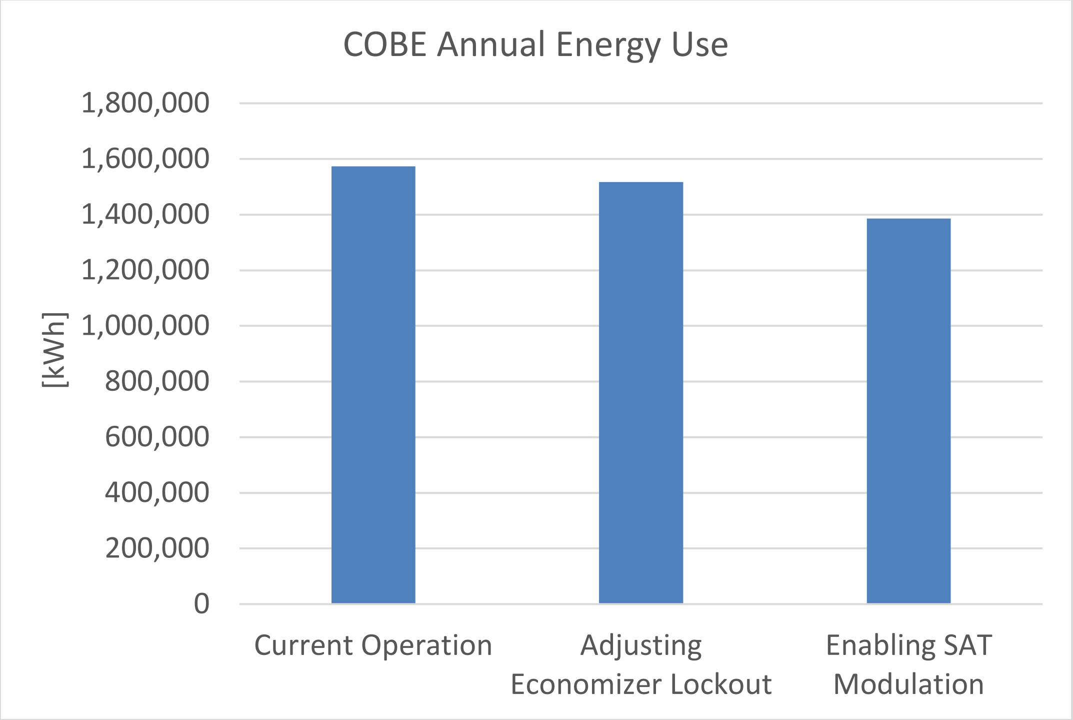

Figure 3.5: Potential site savings identified in the project.

Research Contributions

One of the unique aspects of this research is that it was performed almost entirely off-site. Our research team visited the site twice briefly to verify modeling and operational assumptions, but performed the bulk of the analysis in a remote laboratory. Rather than connecting the energy model directly to the controller at the site, we acquired a duplicate controller of what was in place at the site and loaded the same operational logic onto it. A duplicate controller was easily procured from surplus inventory at the facility, but the research team also had to verify that the same logic was loaded onto the machine, and that we could access to the manufacturer's native program for BACnet to identify each control point as shown in Table 3.1.

During the testing of BCVTB, we found it easiest to send signals through BACnet IP as opposed to Ethernet or MSTP. Field controllers use a mix of signals in their logic.

These signals include inputs, outputs, and values, and can be binary or analog. The BACnet protocol includes 18 different signal or “object” types [67]. Two of the object types that the controller used included analog inputs, which serve as sensor inputs, and analog values, which serve as a setpoint or other analog control system parameter. The objects are typically programmed by the controls engineer before they are set in the field to receive a signal from a specific device such as a thermostat. The Alerton controller did not allow analog inputs to be over-written. Therefore, all variables required for the damper control had to be changed over from analog inputs to analog values. By changing the BACnet object type, we were able to edit these channels to the numerical values we desired. This conversion was done in the controller’s native program, Envision for Backtalk. While the process of replacing analog inputs with analog values only required an hour of time for an experienced user, it was still an inconvenience and presented a minor hurdle to controller-to-model connection and testing for the Alerton controller. The Flexstat and Reliable Controls software did not have this particular barrier, which expedited the co-simulation. These vendor-specific issues are important for evaluating market penetration potential and anticipating routines needed for successful program or technology implementation at a large scale.

We discovered that the default settings in EnergyPlus limit the model’s ability to minimize the damper position. This was due to two causes: strict Indoor Air Quality (IAQ) controls, and the modelling tool assuming ideal operation. EnergyPlus will maintain indoor air quality according to specific building standards, ventilating the building optimally to meet these standards [68]. Once we discovered these settings, we could adjust them, but we were unaware of these defaults at the beginning of this research. Another example of the EnergyPlus model's idealized operation is that the EnergyPlus model will use as much free cooling as possible. This causes the damper to operate in an ideal manner, opening the damper more than either of the physical controllers signaled.

For dry climates, it is advantageous to use the economizer when the outdoor air dry-bulb temperature matches that of the return air temperature. However, this is not the case in humid climates as the outdoor air can carry significant latent loads. Thus, in a humid climate it is best to manage the economizer with enthalpy control as opposed to a dry-bulb temperature lockout. It is possible that the KMC and Alerton controller tuning was conservatively set so that the controller would function well in a range of climate zones – opening the damper less in dry climates than might be ideal. The investigation revealed that adjusting the economizer control based on the climate can significantly impact operation. It also influences how much outdoor air is brought into the building and the control adjustments merit future research. In addition to the tuning adjustments required for the specific climate, this commissioning also revealed how some of the potential benefits of a controller might be under-utilized. The Alerton controller at the site was capable of modulating the supply temperature instead of keeping it static, thus allowing more dynamic operation of the economizer, yet this feature was not enabled at the site. The 12% savings identified as a result of implementing this ability in the controller accounted for the fact that there might be additional fan energy used as a result of modulating the supply temperature.

Another lesson learned in the process was that the energy simulation using the full loop communication cycle took much longer than an energy simulation on its own (without hardware-in-the-loop). On a typical Windows 7 laptop with four gigabytes of RAM, the model simulation required about five minutes to run the model through a full year of weather data. Adding the controller hardware-in-the-loop increased the simulation time substantially. A simulation required about 24 hours to run through one full year of AMY weather data. The sending of each signal through the middle-ware interrupts the normal calculation process within the energy simulation, causing the delay in simulation time. The first connection to the controller significantly slowed the simulation. Adding subsequent connections to the controller, however, slowed the simulation down only incrementally for each additional point.

The EnergyPlus user documentation recommends using at least four timesteps per hour for a traditional simulation, with a minimum of one timestep per hour. However, when running a co-simulation, this guideline does not apply. When we used fewer timesteps we observed more controller signal oscillations. While increasing the size of the timesteps allowed the model to run significantly faster, it certainly reduced the resolution of the results. When the model was running at four time steps per hour, the simulation would complete a year’s run within one day of clock time. The shorter the timesteps however, the better EnergyPlus will be able to handle dynamics and larger timesteps (even at 15 minutes) could be a source of controller error.

The weather file we used in the simulation (like most) reported conditions only on an hourly basis. EnergyPlus performs a time-weighted hourly interpolation of this data, but at its heart, the data is still fairly coarse compared to the typical inputs received by a controller. For example, one of the analog inputs on the Alerton VLCA-1688 for the mixed air temperature was designed to receive data at a sampling rate of many times per second. For a TMY file, the minimum space between steps is one minute. This can be overcome by using a custom AMY file with finer measurement resolutions, but we did not explore this approach in this project. Therefore, when running coarser timesteps, BCVTB provided the controller with a constant signal in between each step. This situation is a departure from a strict co-simulation approach, which syncs the building to the controller in real-time. Instead, the goal here was to approximate building behavior for one year to quickly identify areas of controller tuning. Yet, even this compromise of taking a day or more of clock time to run simulations at fine timestep resolutions is a significant improvement over waiting for an actual year to pass as is required for a best-case scenario of physical commissioning. This virtual commissioning approach can add value to those in the controls industry. The project raised awareness of the potential for simulation-based pre-commissioning. Testing control sequences prior to installation at a site is of great value to industry professionals, as this would mitigate some of their operational risk. Utility companies and building owners would also benefit from the reduction in energy consumption.

Conclusion

We constructed a calibrated EnergyPlus model and established a co-simulation between BCVTB and EnergyPlus that communicated with EMS hardware-in-the-loop. We studied a 50,000 ft^2 three-story building serving a mix of classrooms and offices and modeled it in OpenStudio/EnergyPlus. We set up full-loop communication between this energy model and a duplicate of the physical controller at the building facilitated via BCVTB. Our research provided a practical application of the methods outlined in earlier work from LBNL [54]. We replicated and extended these methods, and proved savings potential from its practical application identifying, overcoming, and documenting several problems that are necessary to repeat the method and support commercialization. This research was performed to help enable incentive programs or other value-added energy services to be developed, which will improve the effectiveness of new building commissioning, existing building retro-commissioning, and promotion of innovative designs for high performance buildings. The project also served as an outreach to two different control companies, acquainting them with simulation-based pre-commissioning. This research showed a direct savings for a physical building and provided the researchers the means to begin expanding the use of this technology as a value-added service. Almost all of this work was performed remotely without the high expense of a day’s worth of time for a commissioning agent and the facilities team. By using a controller that was separate from the actual building at the site, we were able to run full annual simulations and record the controller responses without waiting for specific conditions to occur in real time.

Connecting a calibrated energy model with a physical controller is possible and it provides a streamlined diagnostic tool to identify control faults that might otherwise have gone undetected. When the research team followed up on the project a year later, the economizer was still set to a low lockout temperature. After further investigation, the facilities team found that the outdoor air sensor for that building was getting showered every afternoon by the irrigation system due to its location. This led to a very inaccurate recording of the outdoor air temperature at the site and was the reason why the economizer was set so low. It is unlikely this would have been discovered without the virtual commissioning research. The virtual commissioning process enables a fast comparison between a building’s actual control settings versus the idealized settings in EnergyPlus. Contrasting the two controls allows for a full analysis of the system without monopolizing the time of the building operators. This shows the pathway for a potential service that can test control schemes for a new structure even before it is complete, or provide continuous re-commissioning or retro-commissioning for existing buildings. The co-simulation of a full-year required a great deal of computational time. In order for this to be feasible in a model-predictive control approach, the simulations must be shortened to provide control feedback in real-time. The following chapter provides an overview of how the energy models of radiant systems were set up to run on short time-horizons in a predictive mode to provide real-time control feedback for co-simulation.

CHAPTER 4: USING FORECASTS AND MODELS TO GUIDE THE CONTROLS

Introduction

In order to run the controls based on predicted loads and responses, a framework is needed that can support the simulation inputs and process the outputs. This framework was based on the structure of the Spark tool where parametric simulations can provide rapid feedback on a range of different scenarios. This framework functions as an adaptation of model-predictive control for making control decisions. Traditional MPCs use a simplified equation instead of a full energy model. In this case, the model was used more as a guide – similar to the virtual commissioning process of Chapter 3. Instead of using a full year of weather, the simulation time was shortened to two days at most, and sometimes as little as eight hours of weather inputs. This enabled rapid feedback and provided a window into how the building was about to perform. The simulation would run three scenarios: raising the slab temperature, keeping the slab temperature the same, or lowering the slab temperature. The outputs of the simulation included comfort and energy metrics: the percentage of people dissatisfied, and the source energy consumption per square foot. The first step in the model guided control process was to create a weather file for the simulation composed of forecasted data.

Methods: Compiling the Weather Forecast File

Developing the weather forecast file involved three specific phases:

1. Collecting a weather forecast for a specific time and place.
2. Generating solar radiation forecasts.
3. Formatting the information for EnergyPlus.

Collecting Forecast Data

The first step to running the energy models in predictive mode is to have consistent access to reliable weather forecasts. Weather forecasts can be accessed in any number of ways. Many use data from the National Oceanographic and Atmospheric Association (NOAA). At times, the NOAA data can have gaps or errors in it. This research relied on a third-party weather resource: the DarkSky Application Programing Interface (API). The Dark Sky Company is a self-funded software startup that specializes in local weather forecasts [69]. It uses bots to access a range of weather forecasts, collect NOAA satellite information and filters it all. The service is used by Microsoft, Yelp, ConEdison, and others. Users register for an API key that gives them access to 1,000 free queries each day. The data analyst Bob Rudis built an R-script that leverages this API and uses a j-son wrapper in order to download this data for a given set of coordinates on command [70]. The research team adapted Rudis' script in order to download a two-day forecast for a location in an hourly format. The script then sends this output to a comma delimited file for further operation.

EnergyPlus models require 35 weather data points to perform a complete simulation. The fields of interest from the weather forecast include the timestamp, precipitation intensity, precipitation probability, dry bulb temperature, dewpoint, humidity, atmospheric pressure, wind speed, wind bearing, cloud cover, and visibility. Many of these values could be directly used in the EPW for the simulation, with careful attention to the units used in the EPW as described in the EnergyPlus Auxiliary Programs guide. However, many solar fields required by the EPW had to be calculated. Most of these 35 weather aspects have to do with the quality of sunlight as this has a major impact on daylighting calculations. Very few weather stations provide the level of detail on solar information contained in the EPW. There are several research stations in the Pacific Northwest that provide some of this historical information, but I am unaware of any source that provides solar forecasts to the level of detail found in an EPW input file. This detailed solar data was not available from Dark Sky or even in the NOAA data recorded at the Boise airport. Therefore, research included the development of this solar information using solar altitude equations and regressions based on the forecasted cloud cover, temperature, and humidity.

Generating Solar Information from the Forecast

EnergyPlus parses the solar radiation in the weather file into several different fields. These include, the extraterrestrial horizontal, extraterrestrial direct normal, horizontal infrared sky, global horizontal, direct normal, and diffuse horizontal radiation. The Department of Energy has determined typical values for each of these radiation fields for locations across the country. These values can be found in the Typical Meteorological Year (TMY) file that is standard for many energy simulations. The solar values in the TMY are derived from the National Solar Radiation Database (NSRDB). While the NSRDB is derived from observed data, most of it is modeled [71].

> “Nearly all of the solar data in the original and updated versions of the NSRDB are modeled. The intent of the modeled data is to present hourly solar radiation values that, in the aggregate, possess statistical properties (e.g., means, standard deviations, and cumulative frequency distributions) that are as close as possible to the statistical properties of measured solar data over the period of a month or year. These data do not represent each specific hourly value of solar radiation to the same or equivalent accuracy as the long-term statistics.”

The data that the NRSDB models are based on comes from about 40 stations in the United States. The stations and data for the Pacific Northwest comes from the

University of Oregon Solar Radiation Monitoring Laboratory. The NSRDB models can vary significantly from the historical observations. For the four months of the TMY in Boise, ID when observed data were available [72], the measured diffuse radiation and the modeled diffuse radiation had a correlation coefficient of only 0.67. This is in stark contrast to the observed dry bulb temperatures at the same site which had a correlation coefficient of 0.98 for the same period. The forecasts must rely on some models to predict solar radiation based on forecast sky and temperature conditions. In order to test the effectiveness of the correlations used for developing forecast data, I compared the models to observed data rather than the TMY. This is to avoid comparing a model to a model, and instead only compares the forecast model to observed conditions.

The extraterrestrial horizontal radiation and extraterrestrial direct normal radiation in (Wh/m^2) are a function of the latitude and solar hour. These can be derived using clear-sky solar insolation equations. The extraterrestrial solar radiation is listed as $\boldsymbol{I_o}$ and is the total solar radiation that falls on a spot above the atmosphere. It is based on the eccentricity of the earth's orbit around the sun.

$$I_o = 1367 * E_o$$

$$E_o = 1.00011 + 0.034221 \cos\Gamma + 0.00128 \sin\Gamma + 0.000719 \cos 2\Gamma + 0.000077 \sin 2\Gamma$$

$$\Gamma = 2\pi \frac{n-1}{365}$$

Io = Extraterrestrial direct normal radiation (Wh/m2)
Eo = The Eccentricity of the earth's orbit at a particular time
Γ = The day angle
n = The numerical day of the year

The Horizontal Infrared Radiation Intensity from the Sky measured in (Wh/m^2) is dependent on other weather conditions. For locations where this parameter is not recorded or available (i.e. from forecast data), Walton and Clark have developed an estimation method [73] [74]:

$$Horizontal_{IR} = \epsilon\sigma T^4{}_{drybulb}$$

$Horizontal_{IR}$ = The Horizontal Infrared Radiation Intensity from the sky (W/m^2)
ϵ= The sky emissivity
σ = The Stephan-Boltzman constant 5.6697e-8 (W/m^{2}K^4)
$T_{drybulb}$ = Outdoor dry bulb temperature (K)

$$\epsilon = \left(0.787 + 0.764 ln\left(\frac{T_{dewpoint}}{273}\right)\right)(1 + 0.0224N + 0.0035N^2 + 0.00028N^3)$$

$T_{dewpoint}$= Outdoor dew point temperature (K)
N = Opaque sky cover (in tenths)

The opaque sky cover has a minimum value of 0 and a maximum value of 10. The opaque sky cover is slightly different than the total sky cover. Typically, only the total sky cover is reported for forecasts. The fraction of the cloud cover that reflects the solar radiation is called the opaque sky cover. The opaque sky cover is always less than the total sky cover. The two values are close, and for Boise, ID have a correlation of 0.91. For the forecast estimate, the opaque sky cover was assumed to be equal to the total sky cover.

The other two solar fields of consequence include the direct normal radiation and the diffuse horizontal radiation. These are components of the global horizontal radiation. The global horizontal radiation in (Wh/m^2) is measured hourly at about 40 locations in the U.S. of which Boise, ID is one [71]. Since this radiation data is not contained in the

DarkSky forecast, the global horizontal radiation was generated using the Zhang-Huang model [75].

$$I_h = \left[I_o \cdot \sin(\beta) \cdot \left\{ C_0 + C_1 \cdot \frac{CC}{10} + C_2 \cdot \left(\frac{CC}{10}\right)^2 + C_3 \cdot \left(T_{db_n} - T_{db_{n-3}}\right) + C_4 \phi \right\} - C_5 \right] / k$$

I_o = The extraterrestrial solar radiation (W/m^2)
I_h = Global Horizontal Radiation Intensity (W/m^2)
β = Sun’s altitude (Radians)
CC = Cloud cover in tenths
T_{db} = Outdoor air dry bulb temperature at the current hour
T_{dbn-3} = Outdoor air dry-bulb temperature at three hours previous
ϕ = The relative humidity (%)
C_1, C_2, C_3, C_4, C_5, and k are regression coefficients specific to the location

The estimation was originally developed for locations in China. Therefore, I developed a regression specific to the test site based on historical observations for Boise, ID. The observed global horizontal radiation intensity for all 8,760 hours was plotted against the estimation. Excel solver was used with an evolutionary engine solver for non-smooth data with bounds on the constants near the minimum and maximum observations from Zhang et al. The final regression is shown in Figure 4.1.

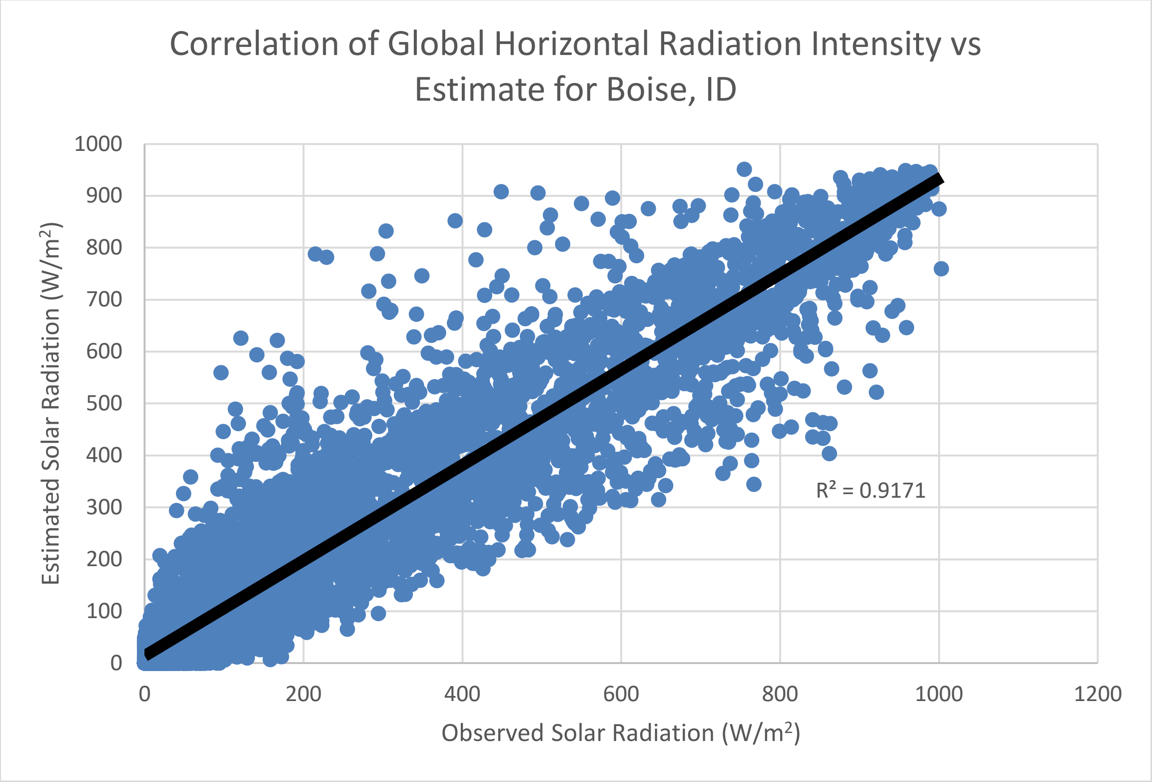

Figure 4.1: Results of correlation between horizontal radiation intensity estimate vs observed data for Boise, ID based on TMY3 file.

The final correlation showed a correlation coefficient of 0.92 and a Root Mean Squared Error (RMSE) of 78. This falls within the correlations developed by the Zhang-Huang model.

Table 4.1: Resulting regression coefficients for Boise, ID from Zhang-Huang model correlation

C_0	0.671205
C_1	0.04125
C_2	-0.32346
C_3	0.004766
C_4	-0.0063
C_5	27
k	0.91427

The global horizontal radiation is further broken down for EnergyPlus into its two components: the direct normal and diffuse horizontal radiation. There are many models available for the decomposition of the global horizontal radiation [76]. The model that showed the best correlation with the recorded data for diffuse horizontal radiation in Boise was the Watanabe model [77]. This is the same model used by Kwak et al. [78]. With the diffuse horizontal and global horizontal known, the Perez model can be used to determine the last remaining solar component [79].

$$Global_{horizontal\ radiation} = Direct_{horizontal\ radiation} + Diffuse_{horizontal\ radiation}$$

$$Direct_{normal\ radiation} = \frac{Direct_{horizontal\ radiation}}{\sin(\beta)}$$

Where β is the solar altitude.

Formatting the Weather Forecasts

Once downloaded and derived, the weather forecast data must be manipulated to fit the file format EnergyPlus models use. EnergyPlus models require a custom file format for weather inputs called an EPW file or (EnergyPlus Weather file). The first seven lines of the EPW contain generic information regarding the location and ground temperatures. After this, each line consists of 35 data points of weather information for one hour. The weather forecast data must be formatted in a very specific manner for it to be compatible with the energy simulation. The run period field in EnergyPlus must be adapted to fit the time-frame of interest. This run period includes the starting and ending month and day for the simulation as well as the name of the weekday on which the simulation is to begin.

While the simulation can be started and stopped for any day, it must start and end at midnight. The EnergyPlus simulation cannot be started at any random hour. As a consequence, the weather data file must have a range of inputs equal to or longer than the run period. For example, if the energy model requests a simulation from 1/25 to 1/27, the EPW must have data for at least 1/25 0:00 – 1/27 24:00. If the weather forecast is from 1/25 07:00 – 1/27 07:00, extra data must be added to the file for the hours of 1/25 0:00 – 1/25 07:00 and from 1/27 07:00 – 24:00. This can be done by either repeating data lines, or stitching together weather history and projections to extend the weather file to the full time required. In this research, the weather forecasts were appended to observed weather conditions to create full data sets for the simulations. The simulation outputs were then filtered for the outputs during the hours of interest. In order to produce this weather data set, the R-script from Rudis was adapted to have a starting and ending date with a for loop that sends an API request for each 24-hour period during that timeframe, appends all of the data, and downloads it to a comma delimited file. The author developed an Excel workbook to automatically parse the forecast data into the EPW order and derive the solar fields based on the equations listed above. The Excel book uses macros to automatically export a comma delimited file in the format of the EPW. Once the file extension is changed it can be used with any EnergyPlus or OpenStudio model.

Methods: Creating the Energy Model

Experimental Site

Figure 4.2: Before and after photos of the warehouse retrofit to CSHQA office courtesy of Ted Isbell [36]

The CSHQA headquarters served as a test-bed for the final analysis. The architecture and engineering office is a 20,000 ft^2, single-story building in Boise, ID. Originally constructed as a warehouse in the 1960s, the building was renovated by CSHQA in 2011. The renovation included adding a radiant slab that could both heat and cool the building. The building is certified as LEED Platinum and won a design award for its innovative HVAC design [80]. The partners at the firm use it as a “living laboratory” to push the envelope of conventional design and practice.

The building plan is open, with a large open office area in the center that takes up most of the space as shown in Figure 4.2. Other spaces in the building include two large conference rooms, a copy room, a sever room, a break area, restrooms, and some private offices. The envelope consists of a brick exterior wall, flat roof with skylights, and large windows along the perimeter. The lighting consists of LED fixtures with integrated photosensors that dim them when enough natural light is present for working. Some of the convective gains from the computer plug loads are alleviated through a unique ducting system. The fans from each computer tower at every working station are ducted into a separate air loop that is exhausted in the summer to minimize the cooling required. Images of the exhaust system are shown in Figure 4.3.

Figure 4.3: The computer exhaust system developed by Russ Pratt [40] [80].

The CSHQA building also has access to a unique resource for heating: the City of Boise’s geothermal system. The hot water available from a fissure in the nearby foothills is available to several buildings in the downtown area and is the largest geothermal network in the United States. Water averaging 160^{o}F is available year-round to those within the network. The hot water is charged by the gallon at roughly the equivalent of what it would cost to use natural gas. CSHQA uses a heat exchanger to extract heat from the geothermal water in order to increase the temperature of its condenser loop.

The HVAC system at CSHQA is a mix of efficient equipment. The system consists of two fluid loops connected through water to water heat pumps. On one side are the individual radiant loops that circulate water through pipes in the floor, on the other side is a condenser loop to which the heat pumps can add or reject heat. The condenser loop of 20% propylene glycol is shown in Figure 4.4. The connection between the condenser loop and the radiant loop through the water to water heat pumps is illustrated in Figure 4.6. The condenser loop maintains a moderate temperature between 60 - 85 ^{o}F.

During the summer, the water-to-water heat pumps (WWHPs) cool the radiant slab by taking the heat absorbed by the radiant loop and rejecting it to the condenser loop. The condenser loop is cooled by a variable-speed, closed-cycle evaporative tower. During the winter, the WWHPs heat the radiant slab by absorbing heat from the condenser loop and pushing it into the radiant loop. The condenser loop maintains its temperature by exchanging heat with the city's geothermal water. An illustration from the OpenStudio model representing the condenser loop is shown in Figure 4.4.

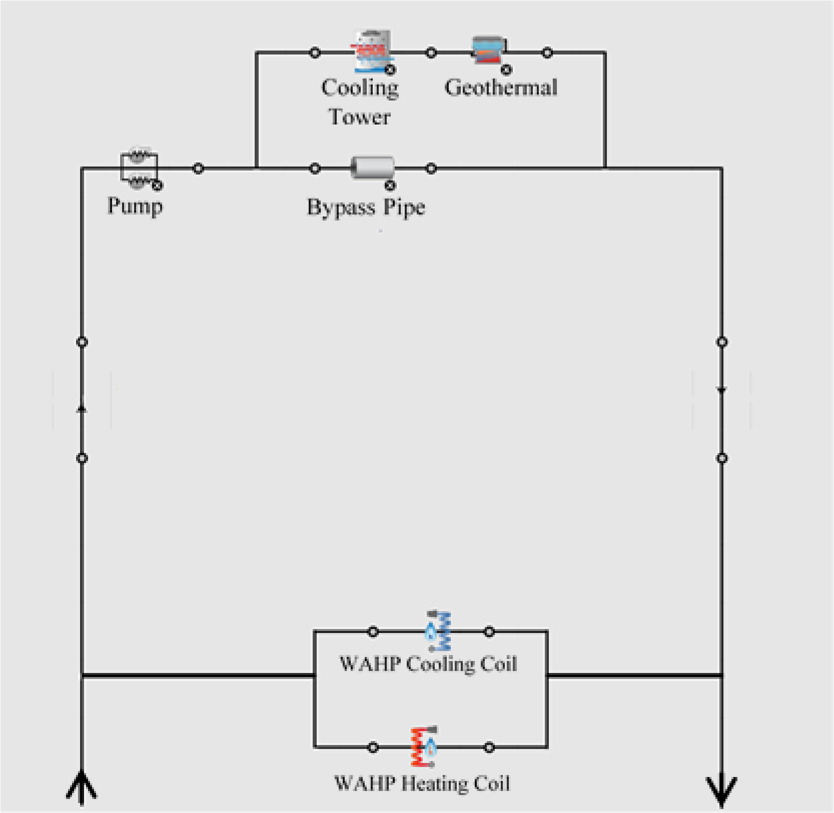

Figure 4.4: An illustration from the OpenStudio software of the condenser loop with a cooling tower and geothermal on the source side serving water to air heat pump coils on the demand side.

The main office area and private offices are designed to be primarily heated and cooled through a radiant slab. Two inches of EPS foam sit beneath six inches of concrete in which 5/8" diameter pipes are set. The pipes are made of cross-linked poly-ethylene (PEX) and there are several circuits of these within the office. There are seven radiant loops. Loop 1 is a 2-ft wide perimeter heating loop that uses water directly from the geothermal system and runs only when the outdoor air temperature is 32^{o}F or colder. Loop 2 serves the four private offices near the east end of the building. Loops 4, 5, 6, and 7 serve the open office area stretching from east to west.

Figure 4.5: An illustration of the CSHQA floorplan from the OpenStudio plug-in for SketchUp.

The four main radiant loops are highlighted in green in Figure 4.5. The radiant system control signal is managed by one central thermostat in the middle of the open area. With the exception of the perimeter loop (shown in red in Figure 4.5), each other radiant loop is heated or cooled through WWHPs attached to the condenser loop. The source side of the coils is the condenser glycol loop and the load side serves the radiant water loop. An illustration of the WWHP layout is provided in Figure 4.6.

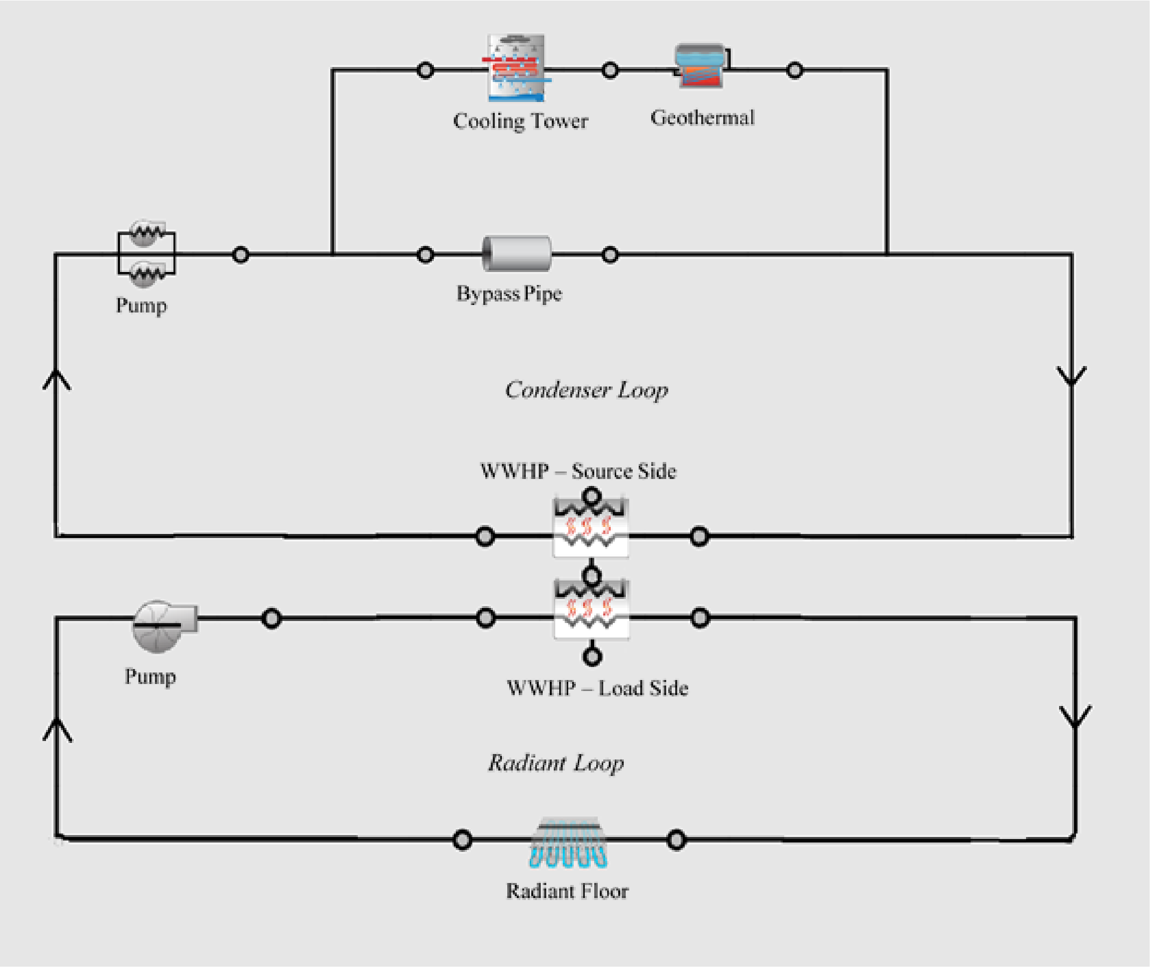

Figure 4.6: An illustration of how the condenser loop interacts with the radiant system as modeled in OpenStudio.

Ventilation and supplemental heating and cooling of the open office area is provided by one main Rooftop Unit (RTU) that is a water-to-air heat pump with a rotary Energy Recovery Ventilator (ERV). The source side of these coils are attached to the main condenser loop along with the WWHP coils. A second water-to-air heat pump serves the copy and server rooms. Three zonal water-to-air heat pumps condition the conference rooms and lobby areas. The restrooms, which include a shower area, have a separate ERV to manage the exhaust and temperature. Small exhaust fans provide ventilation in the unconditioned guest restrooms and garbage space.

Measuring Thermal Comfort Performance at the Site

The research team collected data at the CSHQA site in order to quantify its thermal comfort performance. Rather than rely on occupant surveys or opinions, the team used measured temperatures to compare against standard comfort profiles. With CSHQA's permission, over a dozen instruments were installed in the open office area to collect temperature and humidity readings. These instruments included HOBO-U12 data loggers paired with Type-T thermocouples. Each sensor measured a surface temperature, air temperature, and humidity in the space. The team assumed a default air velocity of 20 feet per minute for a typical work station in the open office area. This enabled the team to estimate comfort based on the thermal comfort standard ASHRAE 55. This standard includes a set of equations for calculating the Predicted Mean Vote (PMV) and the Percentage of Population Dissatisfied (PPD) based on air temperature, mean radiant temperature (MRT), humidity, activity level, air velocity, and clothing level. These calculations are based only on the thermal conditions rather than occupant opinions, which can be highly variable. Instead, comparing the data against the standard allows for an estimation that predicts how the average person would feel in that situation and what percent of the occupants one might expect to be uncomfortable in each circumstance. These equations were formatted into an Excel document so that sub-hourly estimates of the PMV and PPD could be derived based on the data. The data collected in spring of 2017 indicates a low PMV as shown in Figure 4.7 and Figure 4.8.

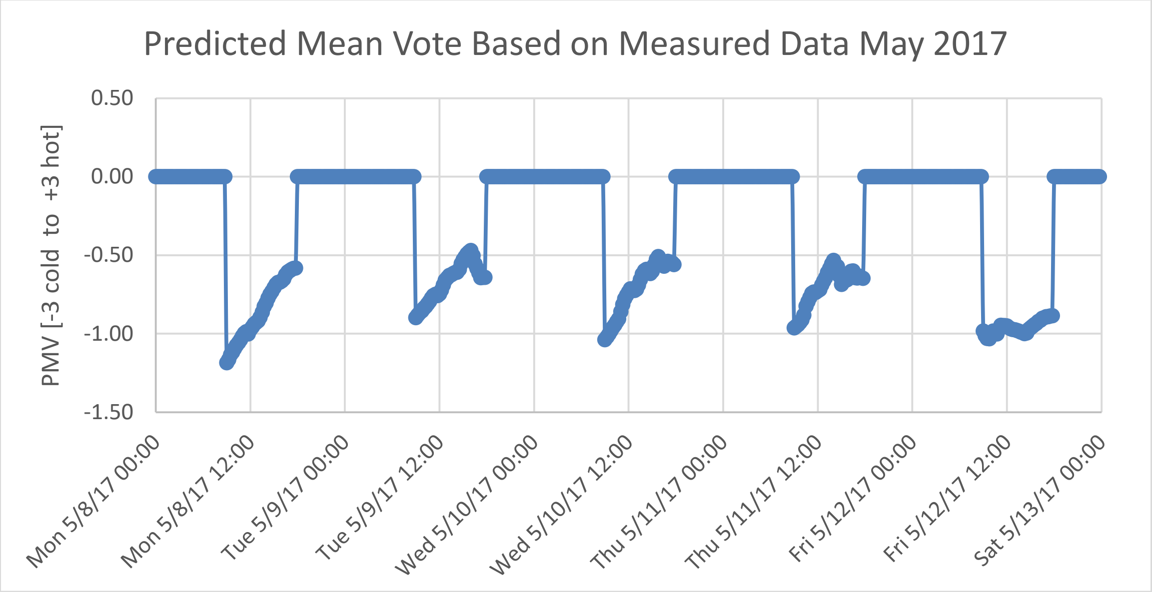

Figure 4.7: Measured thermal comfort performance of the CSHQA open office area in terms of PMV.

The PMV measurements of -0.5 to -2 indicate that the office on average feels rather cold. The thermal comfort standard requires a zone to be between -0.5 to 0.5 for compliance. During this period of measurement, the office rarely achieved a PMV of -0.5 and tended to be colder throughout. The activity level was assumed to be sitting and typing, while the clothing level was assumed to be shoes, socks, trousers, and a long-sleeve button-down shirt based on observations and conversations with the staff. Based on the PMV, one can calculate the estimated percentage of dissatisfied occupants or PPD as shown in Figure 4.8.

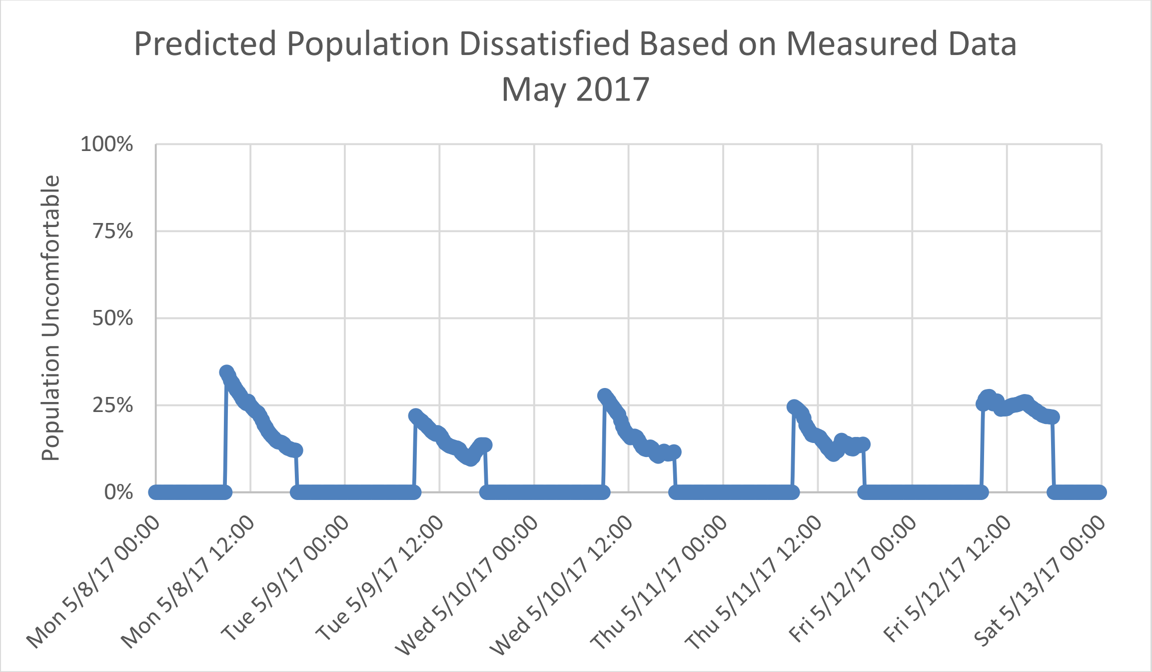

Figure 4.8: Measured thermal comfort performance of CSHQA open office in terms of PPD.

The Predicted Mean Vote is directly related to an estimate of the percentage of people dissatisfied in the space. The data in Figure 4.8 based on temperature and humidity observations show that between 20-30% of occupants were predicted to be uncomfortable working in these conditions. This was a week when the radiant slab was not operating at all and all conditioning for the open office area was provided solely by the rooftop unit. The daily high outdoor air temperature during this period ranged from the mid-70s to the high 80s and the cooling tower was operating during this period. Therefore, the building was in cooling mode even though the thermal performance of the space indicated chilly indoor conditions. The chilly conditions are not unique to this space and several other office spaces have been measured to overcool the spaces beyond what is comfortable. These measurements are part of an ongoing thermal performance study that will conclude in August 2018.

Ideally, an HVAC system should provide comfort and efficiency. If the system does not adequately keep its occupants comfortable, then it fails in its telos. And yet, many systems are controlled based solely on the zone air temperature and do not account for the surface temperatures. A radiant system is ideally suited to target occupant comfort but must be operated carefully to do so. If not, the radiant system may fight with the air side system, failing in both comfort and efficiency. Experimenting with the radiant controls and setpoints can come with a high risk of failure and multiple iterations may be required to determine the best control solution. This is where the model can serve as a guide for the controls.

Calibrating the Energy Model

The energy model reproduced the building's design in exacting detail based on the as-built construction and mechanical drawings. Several site visits and discussions with the designers helped to solidify the details. The geometry was created using the OpenStudio extension in SketchUp. Two models were built using different strategies. The initial energy model was constructed through the software code EnergyPlus 8.2 using Notepad++ to build the text file. The second model was built using the Graphic User Interface (GUI) of OpenStudio 2.5, which relies on EnergyPlus 8.9. Both approaches produced calibrated models and showed that this project could be replicated using either method. The OpenStuido GUI provides many advantages over the text-code only EnergyPlus and adds visual clarity to the model's zoning and HVAC components. Images of the OpenStudio model of the site are shown in Figure 4.4, Figure 4.5, and Figure 4.6. OpenStudio also offers other advantages, including parametric simulations and script-editing measures that make it easy to adjust the model quickly. This is the same functionality the Spark program is built upon.

The model was verified by comparing the simulated performance against measured building performance similar to the co-simulation approach in Chapter 4. The owners of CSHQA provided the research team with access to utility bills that tracked monthly geothermal data and hourly electric data. CSHQA also provided access to the energy management system at the site where trend-logs could be accessed for specific equipment. The building's electrical consumption was sub-metered in the Energy Management System (EMS) to track plug loads, lighting, and HVAC power. The schedules in the EnergyPlus model were adjusted until each of the submetered loads matched the simulated performance of the model. During the fall of 2016, the EMS trend-log feature began to fail and it was no longer possible to track detailed performance.

In April of 2017, a representative from the company BuildingFit [81] installed the program SkySpark on a computer at CSHQA as an alternative way to record EMS data points. SkySpark is a data analytics package that can connect to any BACnet control system. It provides trend-logs and graphics for any point within the control system. ETC Group granted an academic license of 300 points that the research team used to track a range of control signals at CSHQA in real-time. One especially helpful tool was the sub-hourly monitoring of the HVAC power consumption. Having access to the control signals allowed the research team to determine when the slab was active and what the setpoints in the space were. These operational schemes were verified by the controls engineers and commissioners of the system. The weather history for 2017 in Boise was obtained using the forecasting method mentioned above. This allowed a direct comparison of the model to the measured power consumption over the same time period and weather conditions. Figure 4.9 provides an example of a comparison of the HVAC power consumption between the energy model and the actual data logged at the site.

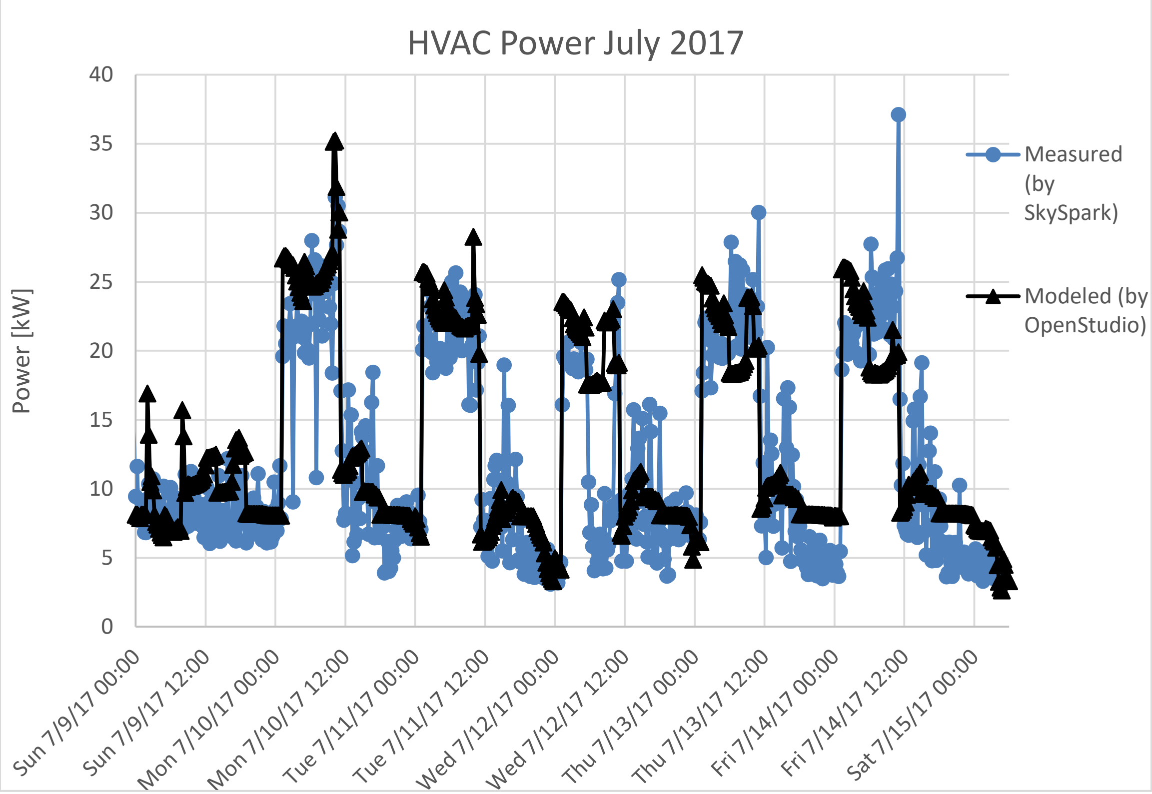

Figure 4.9: Comparison of measured site performance versus energy model simulation.

The model closely follows the actual performance of the building. There are slight variations between the two since the building is subject to dynamic occupancy loads and uses. The model tracks a similar daily consumption with what was measured by the utility as shown in Figure A.3. The only major discrepancies occurred around the July 4th holiday, when fewer workers were present. The energy model was tuned to match not only current operations, but also the indoor temperatures based on data collected by the

SkySpark program during July 2017. The model was able to track this behavior very well as illustrated in Figure 4.10.

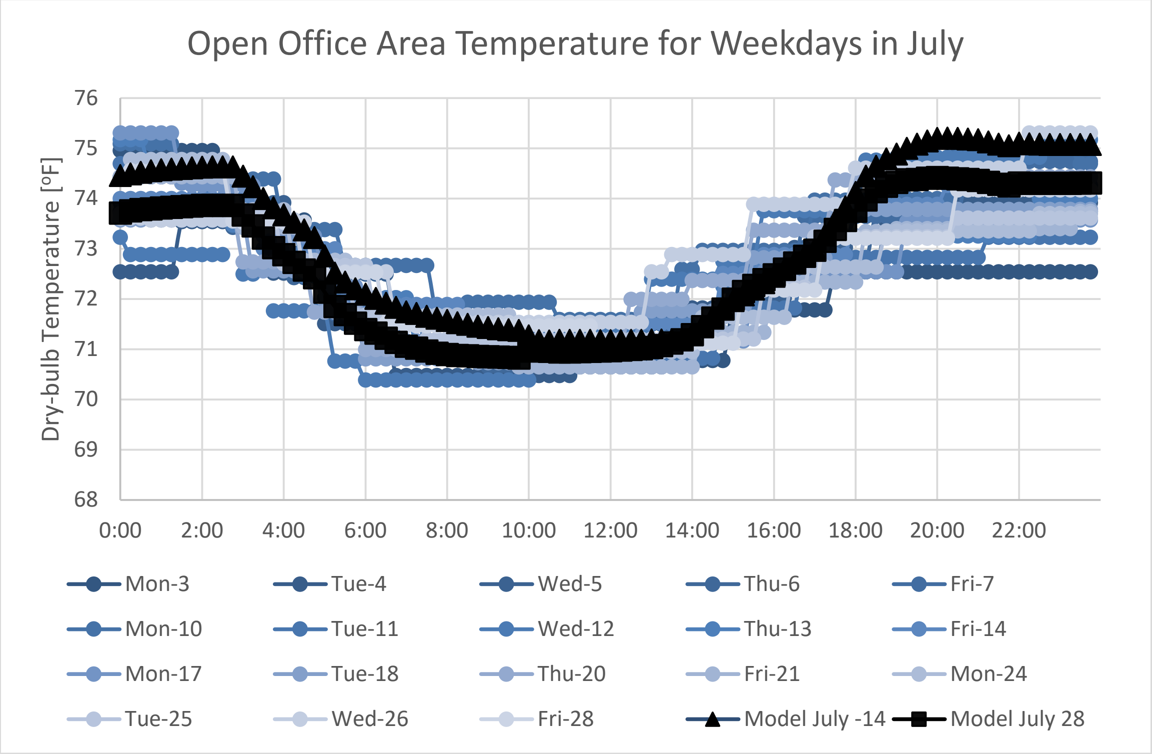

Figure 4.10: Observed thermostat data for all weekdays in July 2017 (Model Outputs are shown in black).

The logged thermostat temperatures from the EMS are listed in blue in Figure 4.10. The stepping of the thermostat signals recorded by the EMS may be due to processing of the signals within the EMS itself as a safety factor to prevent any harm to the HVAC equipment should there be any rapid or irregular readings from the thermostat sensors. This was a feature that the research team noted when connecting a controller to an energy model during the research for Chapter 3. Even though the thermostat signals sent to the EMS appear more stepped than the modeled data, both follow very similar profiles. In addition to modeling the temperatures, the daily HVAC power profile was also analyzed. There is much more variation in the power observations as seen in Figure

4.11 since the HVAC equipment must constantly adjust to outdoor conditions in order to maintain the thermostat setpoints indoors.

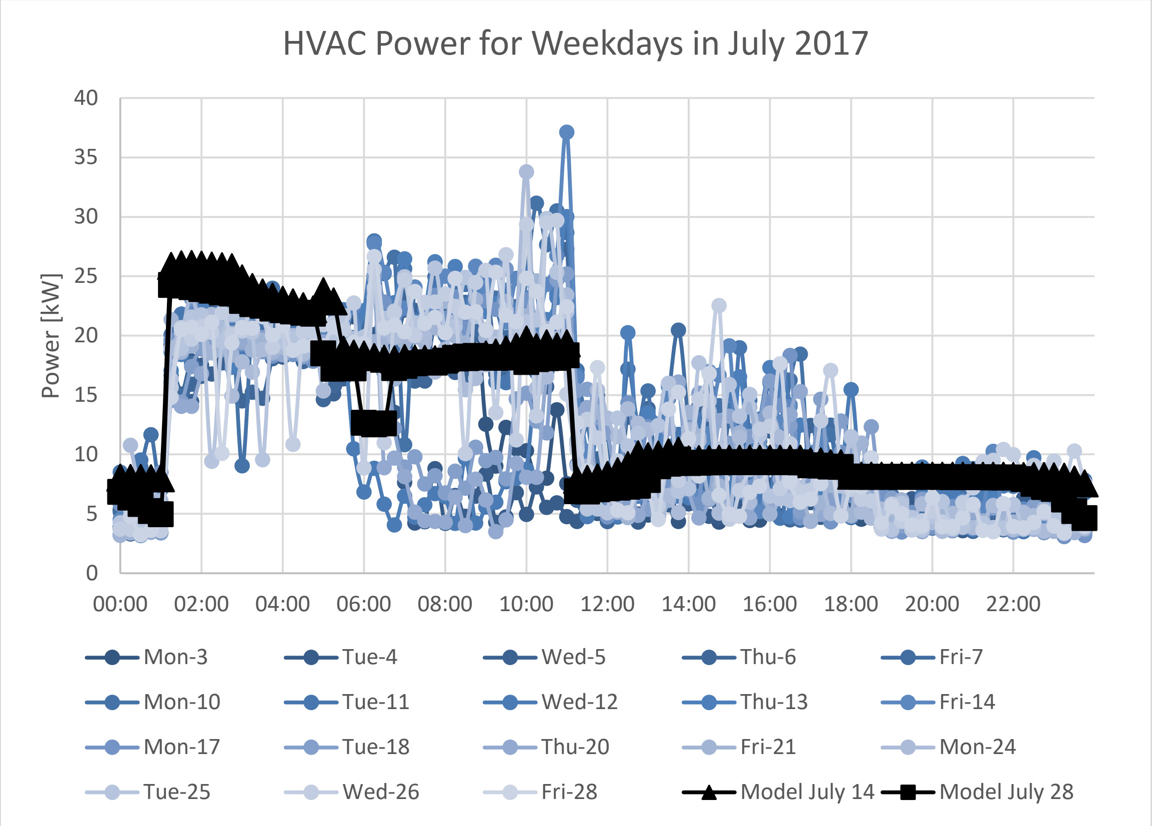

Figure 4.11: Observed HVAC Power Consumption for all weekdays in July 2017 (Model Outputs are shown in black).

The model shared the same average HVAC power profile with the actual use. The actual use (shown in blue in Figure 4.11) had more peaks and variation because some of the equipment at the site cycled on and off, while the model assumed more equipment to be running steadily at part load through things like variable speed drives. While the equipment modeled was the same, the variation came down to actual operation at the site as well as minor daily variations in loading and occupancy. With the model showing reasonable overall performance, it can be used to predict the effects of alternative control schemes including model predictive control.

Methods: Developing the Control Algorithm

All of the tools mentioned thus far were employed to guide the controls to a scheme that could provide better comfort. The radiant system at CSHQA is ideally suited to run a comfortable and efficient operation. Predictive control can rely on weather forecasts. OpenStudio parametric simulations enable the selection of the best control option and Chapter 4 shows the framework for how this signal can be sent to a building's control system. The model was simulated for short periods of time with a range of inputs. After each simulation run, results on the comfort and energy performance are analyzed and the ideal setpoint is selected. This setpoint is used for the next short time period with a new parametric simulation to select a new setpoint and the process is repeated.

Existing controls

The current scheme has the radiant slab limited to only operate between 1:15 AM – 11:15 AM from Monday – Friday. During its operation, the radiant slab attempts to cool the open office air temperature down to 70^{o}F. If the open office temperature drops below 70^{o}F, then it functions as a lockout and the radiant system shuts off. The slab at CSHQA is only active from 1:15 AM – 11:15 AM to avoid interfering with the air-side HVAC. Operating the slab at night also means that the cooling tower will run at night when the outdoor temperatures are cooler and it can lower the condenser loop temperature most efficiently. The main rooftop unit (RTU) operates only between 9:00 AM and 5:00 PM. It typically cycles on and off for only an hour at a time – usually between 11:00 AM and 3:00 PM as the space begins to warm up after the radiant system shuts off. The RTU has a cooling setpoint of 71^{o}F until 3:00 PM. The 18-ton RTU needs to turn on only intermittently to achieve the 71^{o}F setpoint. Driving the setpoint so low causes overcooling of the space during occupancy. Since the radiant system is only available from 1:15 AM – 11:15 AM, this means that the main rooftop unit does most of the cooling. The office setback is at 77 ^{o}F starting at about 3:00 PM allowing the space temperature to slowly rise back up in the evenings when the office is unoccupied.

Control Modification

Once calibrated, the energy model can be used to show the differences between various control scenarios. For example, a change in the setpoints from one model to the next shows how the comfort and energy performance might shift if the same change was applied to the real building. The model can incorporate forecasts by running the simulation at different setpoints. Once the simulation completes, the results can be processed to identify the setpoint that resulted in the best simulated comfort. The first simulation study was performed for the week of July 9, 2017. This was a particularly warm week in Boise, with temperatures routinely reaching above 100 ^{o}F. The existing controls had an air temperature cooling setpoint of 71 ^{o}F for the open office area and the radiant slab was only active between 1:15 AM – 11:15 AM.

The model was run using weather forecasts as a guide to modify the control setpoints within the simulation. This control modification functioned as a rudimentary method to select the best setpoint to provide the best comfort as a modified version of Model Predictive Control (MPC). A traditional Model Predictive Control scheme includes a prediction horizon, a control horizon, and a time step [43]. The prediction horizon is the time into the future that the model looks ahead. The control horizon is the time interval at which a control step will occur. The time step is the time interval at which the model calculations occur. For the first test, the prediction horizon was 24 hours: the model would run a simulation looking ahead one full day. The control horizon was 4 hours: the thermostat control managing the open office area could be adjusted once every 4 hours. The time step was set at 15 minutes: the EnergyPlus model would calculate the heat transfer, comfort calculations, and energy estimates four times every simulated hour. The cost function for this control scheme was split between two values depending on the time. During occupied hours the cost function to be minimized was the estimated discomfort index or Predicted Mean Vote (PMV). During unoccupied hours, the cost function to be minimized was the energy consumption of the HVAC system.

The decision variable was the thermostat setpoint to which the slab would try to condition the open office. This decision variable was reduced to a set of just three values, making it a discrete optimization problem. The three values were to raise the setpoint by 4^{o}F, keep the setpoint the same, or lower the setpoint by 4^{o}F at each four-hour interval. The method for finding the best value of the decision variables was to do an exhaustive search. There were only three possible values that could be selected during each time step. The resulting comfort index (PMV) and energy consumption for each simulation of each control setpoint were tallied in an Excel document. Whichever setpoint produced the lowest discomfort during occupied periods and lowest energy during unoccupied periods was chosen, with priority given to comfort. The horizon was then extended to the next four-hour interval and this process was repeated. A flow chart outlining this process is shown in Figure 4.12.

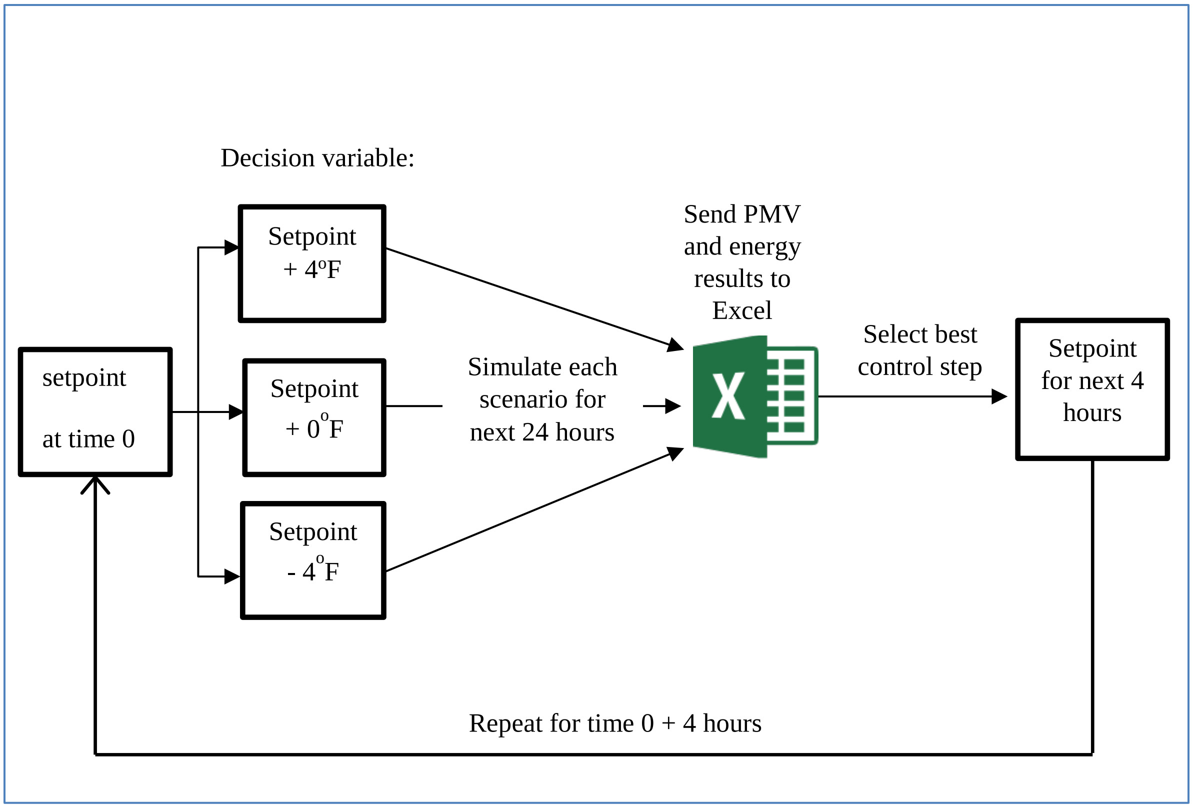

Figure 4.12: Flow chart of the control setpoint selection process.

The control signal for the radiant loop is the thermostat in the open office area. Therefore, this control point was adjusted so that the radiant slab would start cooling at different temperatures. For example, the setpoint temperature was either raised by 4°F, lowered 4°F, or forced to remain the same at each four-hour interval. Whichever response provided the best PMV during occupied hours was chosen as the setpoint for the next four hours. This effectively turned the slab on or off every four hours. The selection criteria in Excel was to simply total the PMV and the HVAC energy for each time-step during the prediction horizon (the next 24 hours). If the control step was to occur during an occupied time, then the scenario with the lowest PMV (best comfort) was selected. If the control step occurred during an unoccupied time, the scenario with the lowest HVAC energy was selected.

Results

In order to isolate the effects of the new control strategy, a calibrated model of the existing controls was compared to a model that used the predictive control method. By comparing one model to another, it removed some of the daily variation in occupancy and loads present in the observed data. When comparing the two control methods, the model guided control scheme did not keep the discomfort index (PMV) perfectly at 0 (see Figure 4.13), but it performed far better than the existing operation and also used 13% less energy (see Figure 4.16).

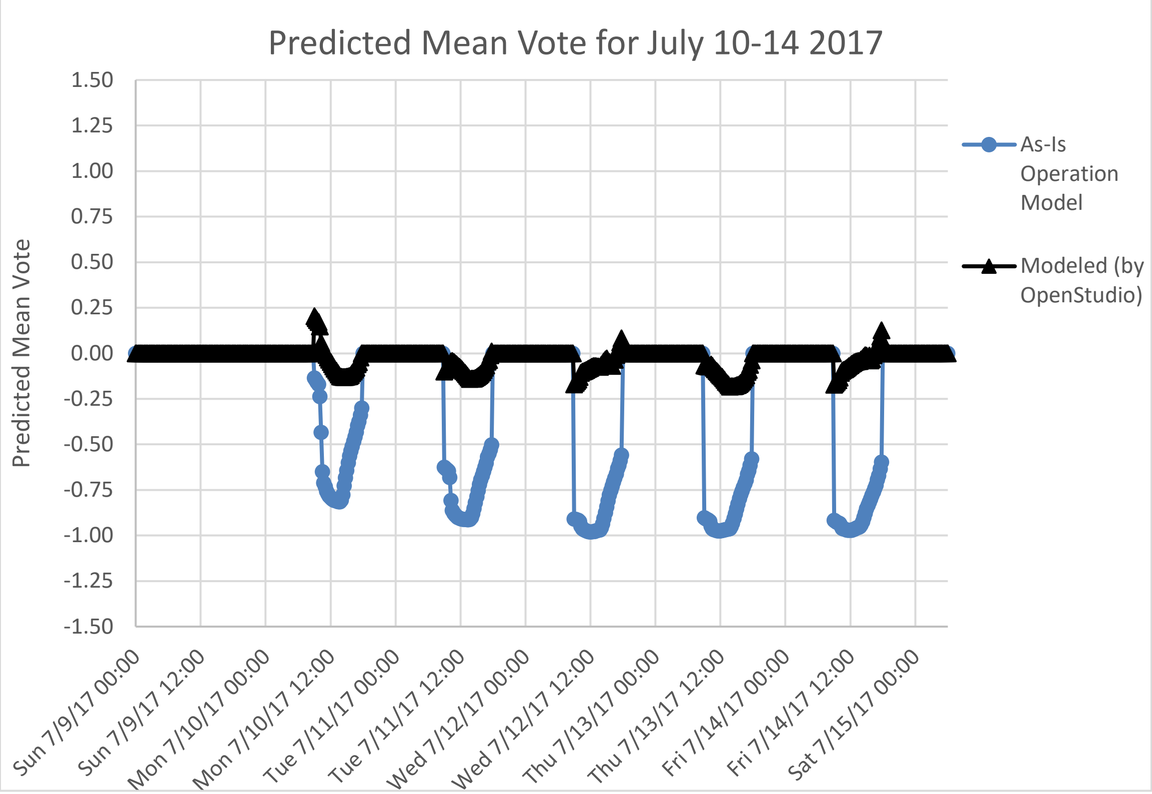

Figure 4.13: Model estimates of performance contrast between existing controls and MPC scheme.

The Predicted Mean Vote for the second week of July showed the MPC scheme was able to maintain a PMV much closer to 0 (optimal) than the existing operation. The PMV is a stand-in for thermal comfort. In this case, the PMV for the existing operation was calculated by the model. The observed PMV from Figure 4.7 showed similar over-cooling trends. By maintaining such a low setpoint, the office is being over-cooled, dropping the PMV below 0. The radiant slab is controlled by the mean air temperature. A contrast between the existing setpoint operation and that developed through MPC is shown in Figure 4.14. One feature of the profile to note is a slight peak in the PMV occurring on Monday morning and again Friday evening. The peak on Monday, where the PMV drifts above 0 is likely due to the control setpoint that was focused on minimizing energy over the weekend (during unoccupied periods) and therefore resulting in a slight penalty for the first control step where the control scheme is focused on minimizing discomfort. In other words, the office is slightly warm after the slab having been inactive for most of the weekend. The control scheme anticipates the occupancy and begins cooling at 4:00 AM on Monday (as seen in Figure 4.14), however, the slight peak indicates the thermal inertia of the slab. Once the slab activation settles into a more regular rhythm between occupied and unoccupied times during the week, the PMV remains closer to 0, although it does drift up again on Friday afternoon when the controls begin selecting for minimal energy as the next interval is during a long period of vacancy.

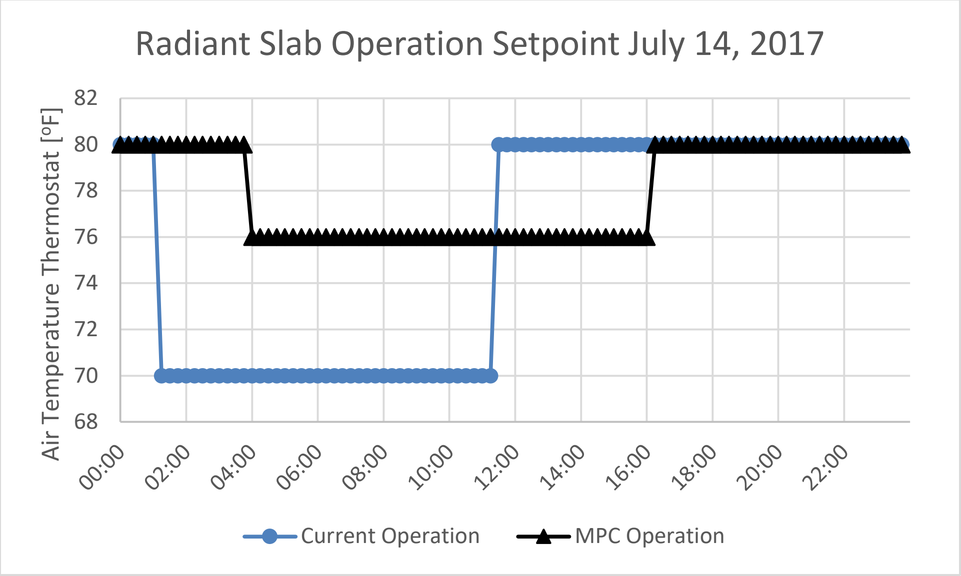

Figure 4.14: Contrasting radiant slab operation versus model guided control setpoint.

The model guided control setpoint achieved better comfort by having the slab cooling start later, operate longer, and maintain a higher setpoint than the current operation. When comparing the slab flow rates as seen in Figure 4.15, one can see that the new control scheme uses the slab in a more varied and extended way than the current operation. The new model guided control scheme used the slab to handle most of the cooling during occupied hours. And, with the new scheme’s higher air temperature setpoint, the RTU became less active and cycled on much less than it had under the existing controls. The overall HVAC power also dropped, indicating that the slab was not negatively impacting the air-side operation. Instead, the air-side system only provided ventilation and back-up cooling as originally intended in the design.

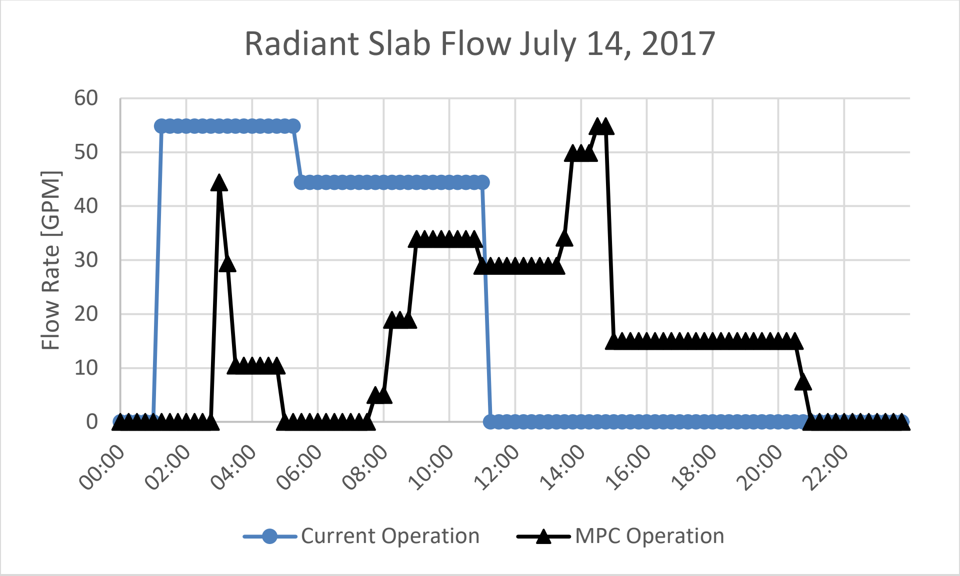

Figure 4.15: Contrasting radiant slab flow from models of current and proposed operation.

The existing operation has such a low setpoint for the slab that the flow is very steady and operates close to its maximum capacity when it is on, before shutting off at 11:00 AM. Since the new setpoint is much higher, the slab no longer needs to operate at its maximum capacity for a brief window of time. Instead, the new setpoint results in the slab cycling on and off more often and running at partial capacity throughout most of the day. The overall temperature in the space is warmer, and the slab would activate more often to prevent the air temperature from rising beyond the setpoint temperature. The flow would pick up as early as 3:00 AM to maintain a setpoint of 80 ^{o}F in the open office area, which had been warming up all evening, and then a small additional amount of slab cooling would be required after 4:00 AM when the setpoint dropped to 76 ^{o}F. The model guided control flow rate shows a later start than the current operation and extends it later in the day. While the overall flow in a day remained similar between the two models, the phase shift of the controls led to lower HVAC energy and a better PMV during occupied periods. What became evident from the model is that there is little need for the use of the rooftop unit if the radiant slab is allowed to operate into the afternoon. The RTU was able to stay off and was only needed to maintain high setback temperatures. A contrast of the HVAC power demand is shown in Figure 4.16.

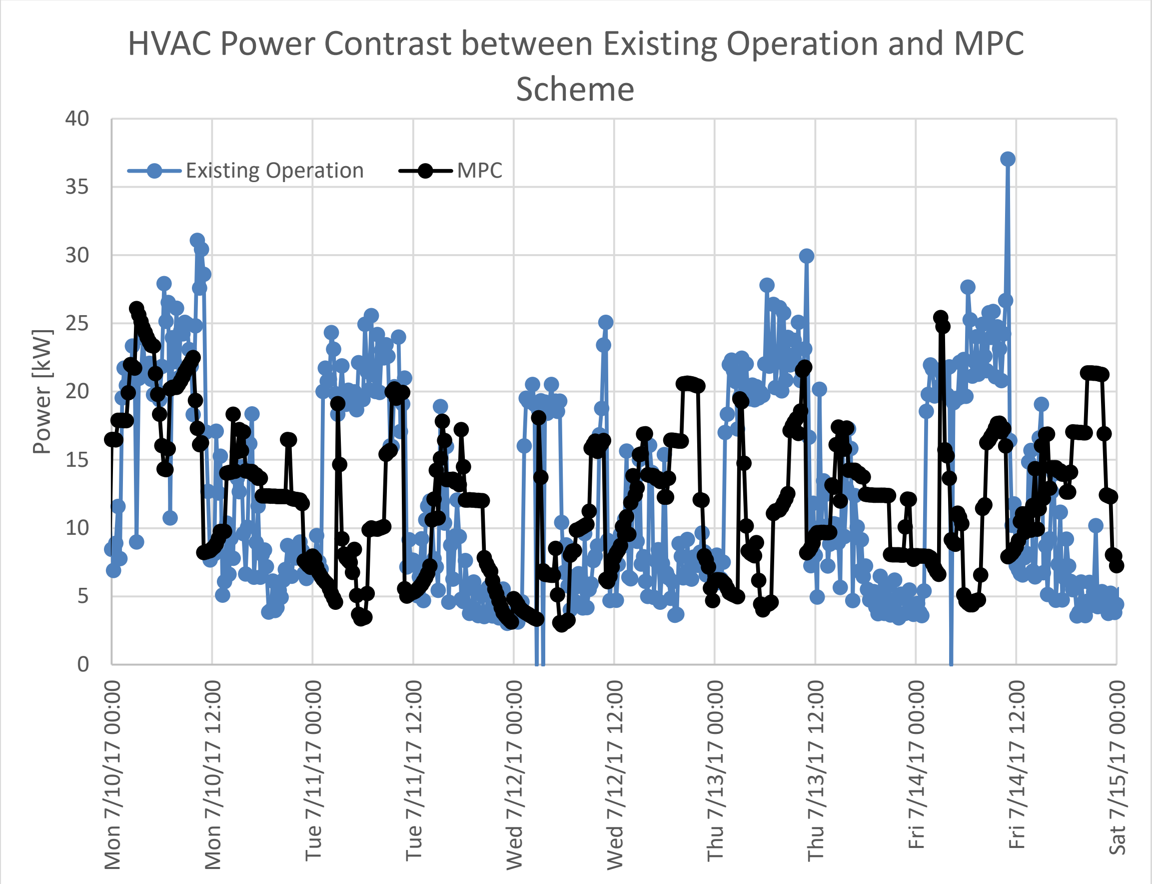

Figure 4.16: HVAC power contrast between observed existing operation and modeled estimate based on MPC.

The MPC scheme also outperformed the current operation in terms of energy. Over the month of July, it was estimated to save about 13% of HVAC energy over the conventional operation. One may also observe from Figure 4.16 that the MPC was able to maintain lower power peaks as well – something of benefit to both the building owner and the utility.

CHAPTER 5: CONCLUSION

Research Overview

This research shows how to reduce energy use and increase productivity in office buildings across the U.S. Taking a holistic view, HVAC is the largest category of use in buildings, which themselves use more energy than any other sector of the U.S. economy. Proper HVAC use is essential to the thermal comfort of employees; that thermal comfort has a direct impact on the productivity of most workers and has a large financial effect. Radiant systems are ideally suited for modern HVAC because these systems can transport heat more efficiently and target human comfort better than traditional all-air systems. Radiant systems benefit from specialized controls to account for their unique dynamics, which are not captured through traditional management by a thermostat.

This study provides a method of managing a radiant slab to provide increased comfort without incurring undue energy costs. This is possible through advanced building controls that include model predictions. Rather than relying on a traditional model predictive control, a full energy model uses parametric simulation results to guide the controls. This research highlights the value of energy models, which are often discarded after the completion of a design. Energy models can be useful throughout a building's life and help to inform operations. Testing at a separate site showed that it is possible to use an energy model to communicate with a control system and even to virtually commission the building controls using a remote co-simulation approach. Therefore, an energy model was built and calibrated for an office in Boise, Idaho that uses a radiant slab to heat and cool an open office area. Weather forecasts were combined with solar correlations to develop forecast files that allowed the energy model to predict the heating and cooling loads in the space for the next 24 hours. Based on these predictions, the slab control setpoint was modified every four hours to select one of three discrete control steps that provided the best comfort during occupied times, and best energy performance during unoccupied times.

When the simulation of the existing controls was compared against the simulation of the model guided control scheme, the results showed that the model guided control scenario saved 13% of the HVAC energy for the month of July. These savings were available because the existing operation maintains the space at such a cold temperature, that it is uncomfortably chilly. Allowing the air temperature to rise and utilizing the radiant system to remove the heat gains results in far more efficient operation. The conventional system limits the radiant slab by cutting off its operation and supplementing the afternoon cooling with the rooftop unit. By using a forecast and a coarse radiant control setpoint adjustment every four hours, the space could be conditioned much more effectively and provide better comfort.

Context and Contribution

This work relies on the contributions of many past research teams in the fields of HVAC, modeling, and control. The final model guided control scheme brings these fields together as part of a new approach to providing efficient comfort in an office space. Past research has profiled the financial cost of inadequate comfort [8] [82]. The research team at the Center for the Built Environment has shown that radiant buildings can provide better comfort than traditional systems [83]. Their team has also shown the inherent efficiencies of radiant systems the benefits Model Predictive Control [11]. The work outlined in this dissertation used some of the equations developed by CBE to profile the comfort in a working office with a radiant slab [21]. An energy model was used to highlight current operations at this site. Unlike a typical calibration, which relies on monthly or hourly total energy consumption [83], this model was calibrated based on sub-metered data from its three major end uses, equipment operation times, and indoor temperature profile through use of the SkySpark data analytics package. While many MPC schemes have been proposed for other building systems, most have tried to minimize either energy consumption or utility bills, while keeping the comfort index within a certain margin [42] [84] [85]. Research around MPC focused on maximizing comfort in the past has been more broad and not specific to radiant systems [86]. The novelty of this control approach was to make comfort a priority and select for the signal that would best serve the occupants of the office specifically for a radiant slab. This research showed that energy savings over the existing operation were possible, even when the primary control motive was to provide comfort – which is the primary purpose of an HVAC system.

Traditional MPCs reduce a model to a set of differential equations so it can be integrated into a control scheme with only a few inputs and outputs. In contrast, this research around model guided controls developed a simplified control scheme around a full energy model. This model uses parametric simulations based on the framework that was developed for a new energy savings estimation program. Several past research teams have used full energy models to anticipate loads through programs like BCVTB. However, most of these studies rely only on outdoor air temperature and humidity as signals of upcoming weather patterns. In contrast, this research relies on holistic weather forecast files based on site-specific solar correlations. Kwak and Seo developed the method for these correlations but applied it only to locations in China and was focused on ice storage [78].

One factor that distinguishes this research from past MPC and co-simulation studies is that the model is decoupled from the control system in this research. In most situations the MPC is integrated into the Energy Management System [63], but this can come with risks to the equipment and may compromise the system's cyber-security. Decoupling the model from the control systems allowed the research team to perform remote virtual commissioning. This decoupled approach gives greater control and autonomy to a building manager who may welcome the idea of using a model as a guide, but be hesitant to give the model control of the entire system.

The current dissertation benefits from contributions to many fields from many past research teams. This project would not have been possible without the recent advances in energy modeling, building controls, and computer technology. These advances have made it feasible to employ energy models in new and creative ways. Using energy models to enhance comfort by guiding the control of a radiant slab with weather forecasts is a novel approach that has practical financial and environmental benefits.

Discussion

The results showed that using model guided control improved comfort and reduced energy use for this site. The energy model can serve as a valuable tool not only during the design phase, but also during a building's operational life. The virtual commissioning of a site through remote co-simulation showed that significant savings were possible if the controls were to be adjusted. However, the research at this site also revealed reasons for using caution when connecting an energy model to a building’s control system. Unfiltered signals from the energy model have the potential to seriously damage equipment. For example, when analyzing the economizer controls, the energy model would send a signal to open or close the damper instantaneously. Were this signal not filtered, this could break the blades of the damper by rotating them too quickly. This is why there is a filter on that signal within the Energy Management System. There are likely many other pieces of equipment which have similar safety mechanisms built into a conventional control system that may not be accounted for within the energy model. This is why it was safer to run the commissioning experiment on a decoupled piece of control hardware when it was disconnected from the building components.

A second danger from integrating a model with controls is that the communication platform between the model and the controller was over the BACnet protocol through BCVTB. The building managers at the site, as well as the controls engineer warned that BACnet is both a read and a write program. In other words, even if the energy model only reads information from the site, the communication protocol itself makes the rest of the building’s EMS vulnerable to cyber-attacks if that model is receiving information from an unsecured location (like collecting weather forecasts).

A third lesson from the virtual commissioning project is that the signals from the EMS may themselves be unreliable, which could lead the model to make poor control decisions. The research team followed up with the building managers at the site to try to understand why the economizer lockout was so low. The team (and the building managers) discovered that the outdoor air temperature sensor that the economizer used was poorly situated in a location that was hit with sprinklers every afternoon. Since the outdoor air sensor was unreliable, increasing the economizer lockout would have sent in very hot air to the building and caused severe comfort problems and wasted energy.

For these three reasons, it was helpful to have the energy model decoupled from the full EMS. These are all issues that could arise from using a traditional MPC that is integrated with the controls. Instead, having a human-in-the-loop provided practical oversite, maintained equipment safety, and did not compromise cyber-security. Yet, by contrasting the model with the current operation, the building managers were still able to commission the control settings and to identify issues with the outdoor air sensor. This model-based commissioning could be used at any site with a calibrated energy model.

The next step of the research used weather forecasts so that the model could guide the controls. Using a full energy model instead of a simplified equation includes many benefits such as the ability to track and report many different variables. However, full implementation of the model into the control system was not considered both for the reasons listed above and because the energy model may have version control and software functionality issues. Over the course of this research project alone, the modeling software underwent close to a dozen different version upgrades. While many of these updates were beneficial, the version control adds another burden to the building manager if it is fully integrated into the system. The Energy Management System itself had several components that failed during the course of this research, like the data logging feature. Were the model to rely on this feature, the whole system could have been shut down. However, having the model decoupled from the controls meant that the system’s resiliency remained unchanged.

Instead of being integrated, the model guided control serves as a reference point for the building manager. The manager may refer to the model's selection of control setpoints should there be an incoming storm that could dramatically change the weather. The current guided control scheme is coarse and only changes its setpoint once every four hours. It is likely that significant improvements could be made to further enhance comfort and reduce energy peaks as discussed in the next section. While using the model to commission a building may occur only occasionally based on historical observations, model guided-control is updated continuously and relies on weather forecasts. Based on the July 2017 simulation results from using model guided control to adjust the setpoint every four hours, the control setpoint varied little during the week. Occasionally, a weather event would alter one of the setpoints, but this occurred only a few times during the month.

Since the guided control scheme was fairly consistent each day, it would be easy for an energy manager to implement new static setpoints and re-examine these each month. The consistent nature of the control setpoints was likely due to the coarseness of the time and control steps. Reducing the control horizon and reducing the temperature control steps are likely to result in a more varied pattern that responds to the weather more dynamically. Yet, even the rudimentary steps still showed that an adjustment of the setpoint profile could result in significant comfort and energy improvements that were feasible to implement. Either model-based commissioning or model guided control could be run by a building manager.

Future Work

This research shows the potential of using weather forecasts and energy models to guide a radiant slab in a way that improves comfort and saves energy over its existing operation. Many improvements could be made on this method to improve the process speed and increase savings. One of the major improvements would be to automate the weather forecast creation. For this project, the gathering of the forecasts was automated using R and Windows Task Scheduler, however, the correlation of solar data and formatting of each forecast file was done manually. Automating the last two pieces of this process would make using this method faster and easier.

After automating the weather forecast creation, the next step to improve this method would be to reduce the control horizon. For example, reducing the control step to every two or three hours. Reducing the control step might have the added benefit of smoothing some of the HVAC power peaks if the change between each step is smaller. The discrete control variable selections could also be expanded beyond the three choices. One issue with using discrete variables and an exhaustive search for the lowest cost is that as the control steps become finer, the number of simulations required increase exponentially. This is one of the disadvantages of the current method and why traditional MPC is so valuable.

There may be a limiting case to reducing the control horizon depending on the characteristic time step of the slab. Another experiment to test would be to use the model to derive that characteristic time step by modifying the thermostat and weather file in the model so that the model produces a step response. This step response could be used to create a reduced-order-model. One could then implement a traditional MPC approach using this reduced-order-model to come up with some control profiles. The control profiles could be tested and compared to the model guided-control approach using the full EnergyPlus model.

Implementation of the model guided control could be further extended and monitored at the CSHQA site. A new round of temperature measurements will be taken during summer 2018 to once again quantify thermal comfort in the space based on the ASHRAE 55 standard. Anecdotal evidence from conversations with those working in the office revealed minor comfort complaints similar to what was measured and graphed in Figure 4.7 and Figure 4.8. However, it would be informative to add official occupant comfort surveys to the data.

The proposed model guided controls are based on continuous weather updates. The model could be used to develop specific responses to more general weather patterns. These responses could be collected into a reference guide for the building operator. For example, if a warm weather pattern is approaching, the building operator could refer to the guide for suggestions for setpoints and scheduling the slab operation.

Comparing the simulation results between the existing controls and the model guided controls showed an energy savings of 13% for the month of July. One way to help verify this estimate would be to implement the model guided controls at the site. This would mean using the model to generate the ideal control setpoints and for the building manager at the site to implement these setpoints over several days. During this time, the open office area should be set up with instruments to estimate the resulting comfort index and compare it to another time using the existing control scheme. Finally, the results should be normalized by taking into account the weather and occupancy differences between the two data sets.

This research could also be extended to include new sites. One of the advantages of using energy models is that each one is specific to a particular building and location. The case study provided in this dissertation sets forth the method and initial findings. Further testing and implementation at sites around the country and around the globe could help many more office buildings achieve better comfort and lower energy costs by using energy models to guide their HVAC controls.

REFERENCES

[1] D. Woods, "Radiant Design Guide," Northwest Energy Efficiency Alliance, Boise, ID, 2016.

[2] D. Woods and J. Gardner, "Using the Time Delay of Radiant Systems to the Grid's Advantage," in *Energy Policy Institute Research Conference*, San Francisco, 2014.

[3] D. Woods, A. Mahic, K. Van Den Wymelenberg, J. Jennings and J. Cole, "SImulation on Demand for Deep Energy Retrofits," in *American Council for an Energy Efficient Economy - Summer Study on Buildings*, Asilomar, 2016.

[4] D. Woods, T. Noble, B. Acker, R. Budwig and K. Van Den Wymelenberg, "Optimizing Economizer Operation by Virtual Commissioning through Remote Co-simulation," in *Building Simulation Conference*, San Francisco, 2017.

[5] Energy Information Administration, "Commercial Building Energy Consumption Survey," Department of Energy, Washington D.C., 2015.

[6] Energy Information Administration, "State Energy Data System," Department of Energy, Washington, D.C., 2015.

[7] S. Cox, Losing Our Cool: Uncomfortable Truths About Our Air-Conditioned World (and Finding New Ways to Get Through the Summer), New York: New Press, 2010.

[8] S. Tom, "Managing Energy and Comfort," *ASHRAE*, no. 6, pp. 18-26, 2008.

[9] J. Feng, F. Bauman and S. Schiavon, "Critical review of water based radiant cooling system design methods," in *13th International Conference on Indoor Air Quality and Climate*, Hong Kong, 2014.

[10] G. Sastry and P. Rumsey, "VAV vs. Radiant Side-by-side Comparison," *ASHRAE JOURNAL*, vol. 56, no. 5, pp. 18-24, 2013.

[11] J. Feng, F. Chuang, F. Borrelli and F. Bauman, "Model Predictive Control of Radiant Slab Systems with Evaporative Cooling Sources," *Energy and Buildings*, vol. 87, pp. 199-210, 2015.

[12] K. Moe, Thermally Active Surfaces in Architecture, New York: Princeton Architectural Press, 2010.

[13] S. Liescheidt, "Radiant HVAC: The how, where & why," *Engineered Systems*, vol. 31, no. 12, pp. 30-32, 2014.

[14] K.-N. Rhee and K. W. Kim, "A 50 year review of basic and applied research in radiant heating and cooling systems for the built environment," *Building and Environment*, vol. 3, no. 91, pp. 166-190, 2015.

[15] K. Roth, D. Wesphalen, J. Dieckmann, S. Hamilton and W. Goetzler, "Energy Consumption Characteristics of Commercial Building HVAC Systems Volume III: Energy Savings Potential," Department of Energy, Washington D.C., 2002.

[16] E. Fabrizio, S. Corgnati, F. Causone and M. Filippi, "Numerical comparison between energy and comfort performances of radiant heating and cooling systems versus air systems," *Hvac and R Research,* vol. 18, no. 4, pp. 692-708, 2012.

[17] Energy Information Administration, "Annual Energy Review 2017," U.S. Energy Information Administration, Washington DC, 2018.

[18] A. Lovins, "The Negawatt Rvolution," Across the Board, Snowmass, CO, 1990.

[19] J. Carter, *Address to the Nation on Energy,* Washington, D.C.: Gerhard Peters and John T. Woolley, 1979.

[20] P. O. Fangar, Thermal Comfort Analysis and Applications in Environmental Engineering, New York: McGraw-Hill Book Company, 1970.

[21] ASHRAE, "ASHRAE, ANSI/ASHRAE Standard 55-2013. Thermal Environmental Conditions for Human Occupancy," American Society of Heating Refrigerating, and Air-Conditioning Engineers, Inc., Atlanta, 2013.

[22] Washington Administrative Code commercial buildings WAC 51-11C140360.

[23] New Buildings Institute, "Getting to Zero 2012 Status Update: A First Look at the Costs and Features of Zero Energy Commercial Buildings," NBI, 2012.

[24] Center for the Built Environment, "Radiant Systems Research," Lawrence Berkeley National Lab, 13 5 2014. [Online]. Available: https://www.cbe.berkeley.edu/research/radiant-systems.htm. [Accessed 15 January 2018].

[25] T. Isbell, "Our Living Laboratory – The CSHQA Platinum Story," in *Idaho Energy and Green Buildings Conference*, Boise, 2016.

[26] D. Watson and K. Chapman, Radiant Heating and Cooling Handbook, London: McGraw-Hill Education Europe, 2008.

[27] S. S. a. F. B. Jingjuan Feng, "Cooling load differences between radiant and air systems," *Energy and Buildings,* vol. 65, pp. 310-321, 2013.

[28] G. McDonell, "High - Performance Buildings Through Integrated Design," *Heating/Piping/Air Conditioning Engineering. 79 (2).*, vol. 79, no. 2, 2007.

[29] ASHRAE, Handbook of Fundamentals, Atlanta: ASHRAE, 2017.

[30] B. Oleson, "New European standards for design, dimensioning and testing embedded radiant heating and cooling systems Clima 2007 WellBeing Indoors," in *Clima 2007 WellBeing Indoors,* Helsinki, 2007.

[31] C. Stetiu, "Energy and peak power savings potential of radiant cooling systems in US commercial buildings," *Energy and Buildings,* vol. 30, no. 2, pp. 127-138, 1999.

[32] S. Jing, Qingyun and He, "Research on suitable application sites of Radiant cooling system," *Energy Procedia,* vol. 14, no. 2, pp. 1002-1007, 2012.

[33] B. Olesen, K. Sommer and B. Duchting, "Control of slab heating and cooling systems studied by dynamic computer simulations," *ASHRAE Transactions,* vol. 108, pp. 698-707, 2002.

[34] P. Weilzmann, Modeling Building Integrated Heating and Cooling Systems, Copenhagen: Department of Civil Engineering, Danmark Technisk Universite, 2004.

[35] B. Olesen, "Possibilities and Limitations of Radiant Cooling," *ASHRAE Transactions,* vol. 1, no. 7, p. 103, 1997.

[36] J.-H. Lim, J.-H. Jo, Y.-Y. Kim and K.-W. Kim, "Application of control methods for radiant floor cooling systems in residential buildings," *Building and Environment,* vol. 41, no. 1, pp. 60-73, 2006.

[37] G. Rumsey and P. Sastry, "VAV vs. Radiant Side-by-Side Comparison," *ASHRAE,* vol. 56, no. 5, 2014.

[38] F. C. F. B. a. F. B. Jingjuan Feng, "Model predictive control of radiant slab systems with evaporative cooling sources," *Energy & Buildings,* vol. 87, pp. 199-210., 2015.

[39] L. A. a. M. Alves, "Uncompromising and Occupant-Centered: Changing How We Engineer Buildings," in *ACEEE Summer Study*, Asilomar, 2016.

[40] T. Isbell, " Our Living Laboratory – The CSHQA Platinum Story," in *Idaho Energy and Green Buildings Conference,* Boise, 2016.

[41] T. Schelling, Micromotives and Macrobehavior, New York: WW Nortan & Company, 2006.

[42] J. Ma, J. Qin, T. Salsbury and P. Xu, "Demand Reduction in Building Energy Systems Based on Economic Model Predictive Control," *Chemical Engineering Science,* vol. 67, pp. 92-100, 2012.

[43] K. Powell, W. Cole, U. Ekarika and T. Edgar, "Optimal Chiller Loading in a District Cooling System with Thermal Energy Storage," *Energy,* vol. 50, pp. 445-453, 2013.

[44] A. Afram and F. Janabi-Sharifi, "Theory and Applications of HVAC Control Systems – A Review of Model Predictive Control (MPC)," *Building and Environment,* vol. 16, no. 11, 2013.

[45] Northwest Energy Efficiency Alliance, "Spark," BetterBricks, 2017. [Online]. Available: https://buildingrenewal.org.

[46] S. Lee, T. Hong, M. Pietter and S. Taylor-Lange, "Energy Retrofit Analysis Toolkits for Commercial Buildings," *Energy,* vol. 6, no. 112, pp. 1087-1100, 2015.

[47] T. Hong, Y. Chen, S. Lee, S. Tylor-Lange, Z. R., K. Sun and P. Price, "Commercial Building Energy Saver: An Energy Retrofit Analysis Toolkit," *Applied Energy,* vol. 159, no. 12, pp. 289-309, 2015.

[48] Cadmus, "Northwest Commercial Building Stock Assessment.," Cadmus Group Inc., Waltham, 2009.

[49] C. Duarte, K. VanDenWymelenberg and C. Reiger, "Revealing Occupancy Patterns in Office Buildings through the Use of Annual Occupancy Sensor Data," *Energy and Buildings,* vol. 67, pp. 587-595, 2013.

[50] NBI (New Buildings Institute), "Deep Energy Savings in Existing Buildings Case Studies," New Buildings Institute, Portland, 2011.

[51] G. Liu, B. Liu, W. Wang, R. Athalye, D. Moser, E. Crowe, N. Bengtson, M. Effinger, L. Webster and M. Hatten, "Advanced Energy Retrofit Guide: Office Buildings," Pacific Northwest National Laboratory, Richland, WA, 2011.

[52] B. Dong, O. Zheng, Z. Li, D. Luo, M. Shashanka, S. Ahuja and T. Bailey, "An Integrated Infrastructure for Real-Time Building Energy Modeling and Fault Detection and Diagnostics," in *Fifth National Conference of IBPSA-USA*, Madison, 2012.

[53] L. Harmer and G. Henze, "Using Calibrated Energy Models for Building Commissioning and Load Prediction," *Energy and Buildings,* vol. 92, no. 4, pp. 204-215, 2015.

[54] X. Pang, M. Wetter, P. Bhattacharya and P. Haves, "A Framework for Simulation-Based Real-Time Whole Building Performance Assessment," *Building and Environment,* vol. 54, pp. 100-108, 2012.

[55] Z. Li, G. Augenbroe, B. Dong and Z. O'Neill, "Database Supported BACnet Data Acquisition System for Building Energy Diagnostics," Energy Systems Laboratory, College Station, 2012.

[56] L. Wang, S. Greenberg, J. Fiegel, A. Bubalcava, S. Earni, X. Pang and J. Hernandez-Maldonado, "Monitoring-based HVAC Commissioning of an Existing Office Building for Energy Efficiency," *Applied Energy*, vol. 102, no. 2, pp. 1382-1390, 2013.

[57] D. Westphalen and S. Koszalinski, "Energy Consumption Characteristics of Commercial Building HVAC Systems Volume II: Thermal Distribution, Auxiliary Equipment, and Ventilation," U.S. Department of Energy, Cambridge, 1999.

[58] E. Mills, "Building Commissioning: A Golden Opportunity for Reducing Energy Costs and Greenhouse Gas Emissions," *Energy Efficiency,* vol. 4, no. 2, pp. 145-173, 2011.

[59] M. Wetter, "Co-Simulation of Building Energy and Control Systems with the Building Controls Virtual Test Bed," *Journal of Building Performance Simulation,* vol. 4, no. 3, pp. 185-203, 2011.

[60] P. Haves and P. Xu, "The Building Controls Virtual Test Bed - A Simulation Environment for Developing and Testing Control Algorithms Strategies and Systems," in *International Building Performance Simulation Association*, Beijing, 2007.

[61] J. Wen, "Rapid-Prototyping Control Implementation using the Building Controls Virtual Test Bed," Philips, Briarcliff Manor, 2011.

[62] S. Karg, "SourceForge," 15 October 2013. [Online]. Available: http://bacnet.sourceforge.net/. [Accessed 12 November 2015].

[63] T. Nouidui, M. Wetter, Z. Li, P. X, P. Bhattacharya and P. Haves, "BacNet and Analog/Digital Interfaces of the Building Controls Virtual Testbed," in *Conference of International Building Performance Simulation Association*, Sydney, 2011.

[64] X. Pang, B. P, Z. O'Neill, P. Haves, M. Wetter and T. Bailey, "Real-Time Building Energy Simulation Using EnergyPlus and the Building Controls Virtual Test Bed," in *12th Conference of International Building Performance Simulation Association,* Sydney, 2011.

[65] T. Nouidui, "BCVTB forum. Retrieved from Google Groups:," Google, [Online]. Available: https://groups.google.com/forum/#!forum/bcvtb. [Accessed 29 May 2015].

[66] M. Wetter and T. Nouidui, "Building Contols Virtual Test Bed User Manual Version 1.5.0," Center for the Built Environment, Berkeley, 2015.

[67] B. Swan, "The Language of BACnet-Objects Properties and Services," *Engineered Systems*, vol. 13, no. 7, pp. 24-32, 1996.

[68] ASHRAE, "Standard 62.1-2016 Ventilation for Acceptable Indoor Air Quality," American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, 2016.

[69] Dark Sky Company LLC, "Dark Sky API," June 2012. [Online]. Available: https://darksky.net. [Accessed 15 January 2017].

[70] B. Rudis, "R interface to the Dark Sky API," GitHub, 22 September 2016. [Online]. Available: https://github.com/hrbrmstr/darksky. [Accessed 15 January 2017].

[71] S. Wilcox, "National Solar Radiation Database 1991-2010 Update: Users Manual," National Renewable Energy Laboratory, Golden, Colorado, 2012.

[72] University of Oregon, "University of Oregon Solar Radiation Monitoring Laboratory," 12 July 2013. [Online]. Available: http://solardat.uoregon.edu.

[73] G. N. Walton, "Thermal Analysis Research Program Reference Manual," National Bureau of Standards, 1983.

[74] G. Clark and C. Allen, "The Estimation fo Atmospheric Radiation for Clear and Cloudy Skies," in *National Passive Solar Conference*, 1978.

[75] Q. Zhang, J. Huang, Y. Hongxing and L. Chengzhi, "Development of Models to Estimate Solar Radiation for Chinese Locations," *Journal of Asian Architecture and Building Engineering*, vol. 41, 2003.

[76] S. A. M. Maleki, H. Hizam and C. Gomes, "Estimation of Hourly, Daily and Monthly Global Solar Radiation on Inclined Surfaces: Models Re-Visited," *Energies*, vol. 10, no. 134, 2017.

[77] T. Watanabe and e. al., "Procedures for separating direct and diffuse insolation on a horizontal surface and prediction of insolation on tilted surfaces," *Transactions, Architectural Institute of Japan,* no. 330, 1983.

[78] Y. Kwak, D. Seo, C. Jang and J.-H. Huh, "Feasibility study on a novel methodology for short-term real-time energy demand prediction using weather forecasting data," *Energy and Buildings,* vol. 57, pp. 250-260, 2013.

[79] R. Perez, R. Stewart, R. Seals and TedGuertin, "The Development and Verification of the Perez Diffuse Radiation Model," Albany, NY, 1988.

[80] R. Pratt, "2017 ASHRAE technology award case studies: Warehouse Transformation," *ASHRAE*, vol. 59, no. 7, pp. 38-42, 2017.

[81] Building FIt, *Subsidiary of ETC Group.*

[82] O. Seppanen, W. Fisk and D. Faulkner, "Control of temperature for health and productivity in offices," in *ASHRAE Transactions,* 2004.

[83] C. Karmann, S. Shiavon and F. Bauman, "Thermal comfort in buildings using radiant fs. all-air systems: A critical review," *Building and Environment,* 29 October 2016.

[84] ASHRAE, "Guideline 14," Atlanta, GA, 2014.

[85] M. e. a. Avci, "Model predictive HVAC load control in buildings using real-time electricity pricing," *Energy and Buildings,* vol. 60, pp. 199-209, 2013.

[86] P. May-Ostendorp, G. Henze, B. Rajagopalan and D. Kalz, "Experimental investigation of model predictive control-based rules for a radiantly cooled office," *HVAC & R Research,* vol. 19, pp. 602-615, 2013.

[87] J. Alvarez and e. al., "Optimizing building comfort temperature regulation via model predictive control," *Energy and Buildings,* vol. 57, pp. 361-372, 2013.

[88] U.S. Green Building Council, "LEED Rating Systems," Leadership in Energy Efficient Design, 2012. [Online]. Available: https://new.usgbc.org/leed/rating-systems. [Accessed 23 July 2012].

[89] J. Feng, Design and Control of Hydronic Radiant Cooling Systems, Berkeley: UC Berkeley, 2014.

[90] Z. T. a. J. Love, "Energy performance optimization of radiant slab cooling using building simulation and field measurements," *Energy and Buildings,* vol. 41, no. 3, pp. 320-330, 2009.

[91] P. H. a. P. Xu, "The Building Controls Virtual Test Bed - A Simulation Environment for Developing and Testing Control Algorithms Strategies and Systems," in *International Building Performance Simulation Association*, Beijing, 2007.

[92] EIA, "Commercial Building Energy COnsumption Survey - E1. Major fuel consumption by end use," U.S. Energy Information Administration, Washington, D.C., 2012.

[93] NBI (New Buildings Institute), "Retrofit Savings Estimator," 2014. [Online]. Available: http://newbuildings.org/resource/retrofit-savings-estimator/.

APPENDIX A

Comparison of Model Outputs to Measured Data

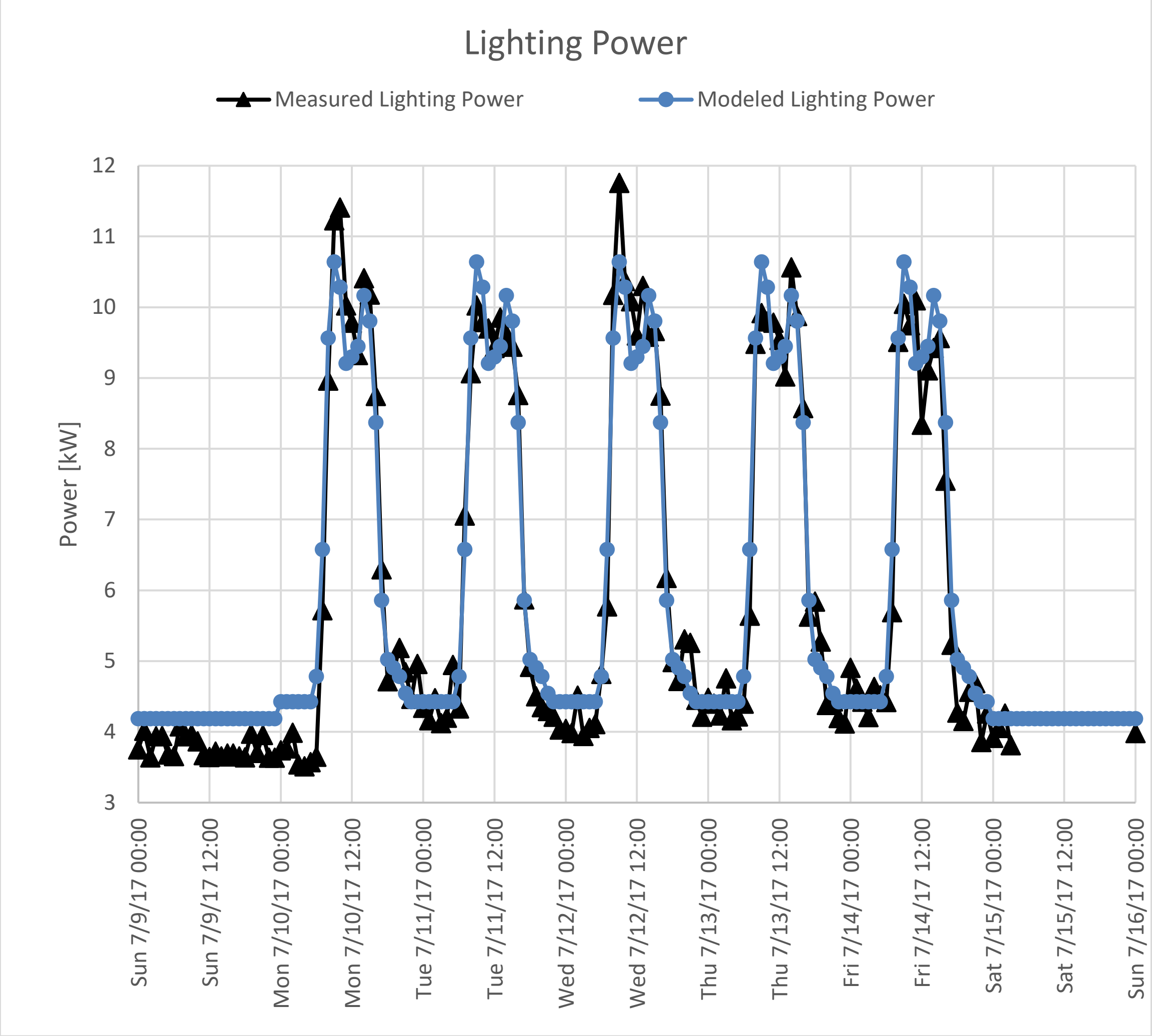

Figure A.1: Lighting power profile for week of July 9, 2017.

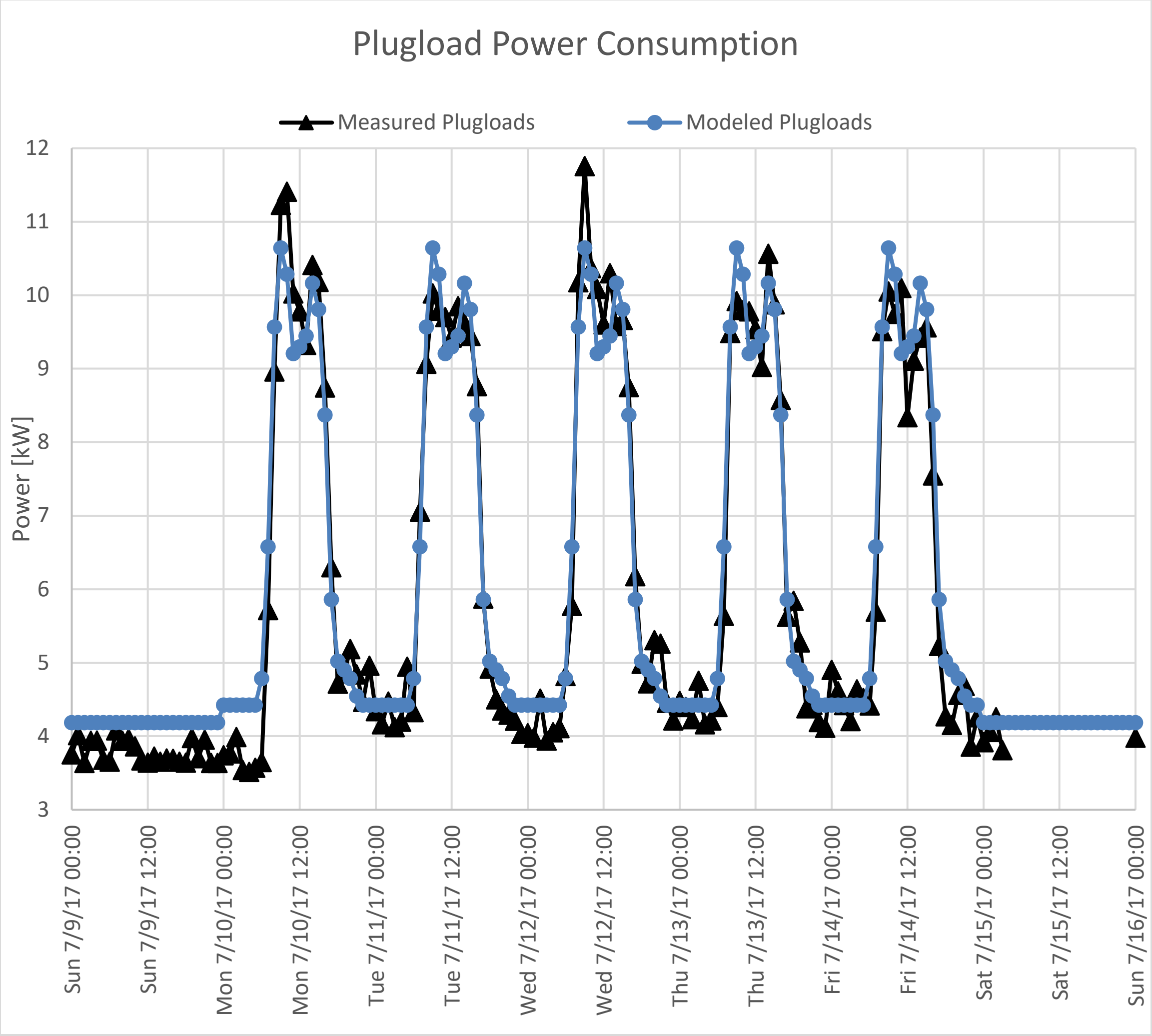

Figure A.2: Plugload power profile for week of July 9, 2017.

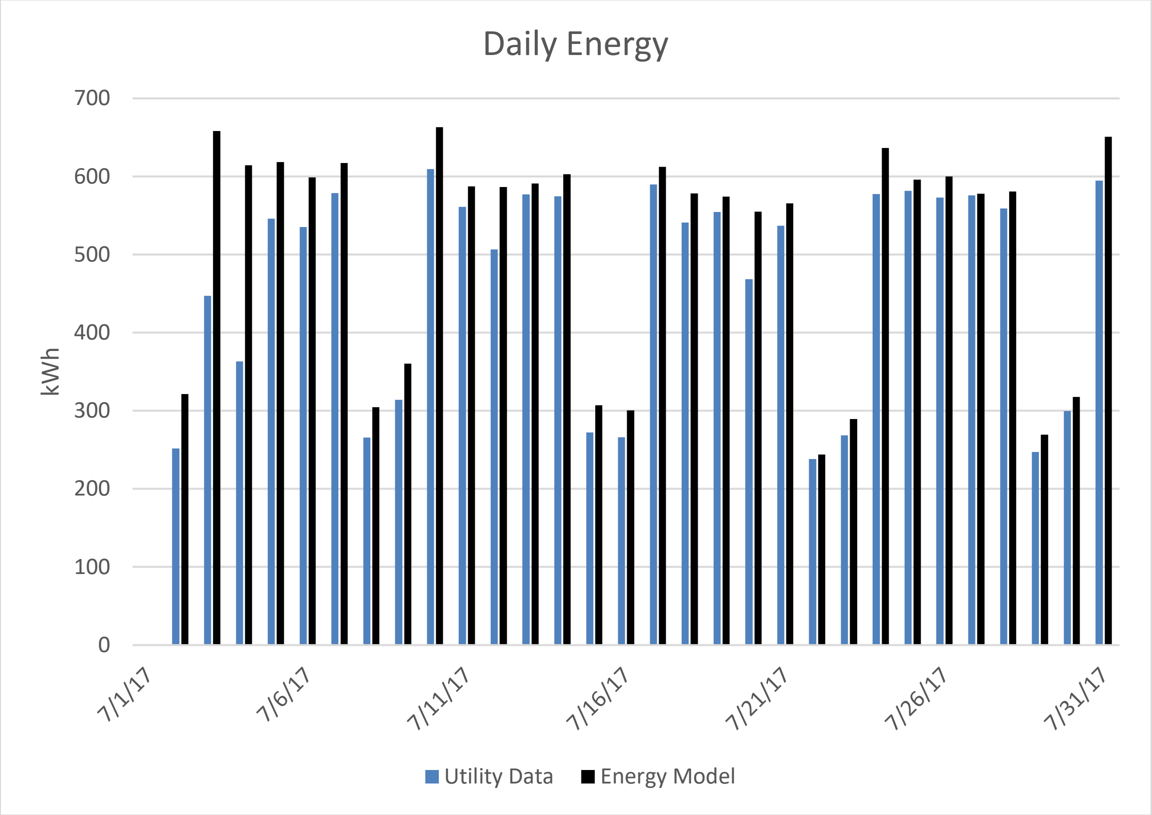

Figure A.3: A comparison of the overall daily energy between the calibrated energy model and the consumption metered by the utility.