

Simulation- and Experiment-Based Setpoint Control for Heating, Ventilation, and Air-Conditioning Systems: A Single- and Multi-Objective Optimization Problem

by

Yuan Cai

B.S. Architectural Engineering
University of Texas at Austin (2019)

Submitted to the Department of Architecture and Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Master of Science in Building Technology and
Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Department of Architecture
Department of Electrical Engineering and Computer Science
January 14, 2022

Certified by
Leslie K. Norford
Professor of Building Technology
Thesis Supervisor

Certified by
Steven B. Leeb
Professor of Electrical Engineering and Computer Science and Mechanical Engineering
Thesis Reader

Accepted by
Leslie K. Norford
Professor of Building Technology
Chair of the Department Committee on Graduate Students

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Thesis Committee

Leslie K. Norford, PhD

Thesis Supervisor

Professor of Building Technology

Department of Architecture

Steven B. Leeb, PhD

Thesis Reader

Professor of Electrical Engineering and Computer Science and Mechanical Engineering

Department of Electrical Engineering and Computer Science, Department of Mechanical Engineering

Simulation- and Experiment-Based Setpoint Control for Heating, Ventilation, and Air-Conditioning Systems: A Single- and Multi-Objective Optimization Problem

by
Yuan Cai

Submitted to the Department of Architecture and Department of Electrical Engineering
and Computer Science
on January 14, 2022, in partial fulfillment of the
requirements for the degrees of
Master of Science in Building Technology and
Master of Science in Electrical Engineering and Computer Science

Abstract

Buildings and building construction sectors are together responsible for 40% of global energy consumption, 70% of electricity consumption, and 40% of carbon emission. Heating, ventilation, and air-conditioning systems (HVAC) in residential and commercial units account for 40% to 60% of energy usage. To increase energy efficiency and reduce energy usage, buildings are now better insulated, installed with energy-efficient appliances, and controlled by advanced technologies to provide user comfort while minimizing their environmental impact.

This thesis focuses on utilizing setpoint control methods to design algorithms for operating thermostatically controlled appliances such as HVAC to achieve the goal of minimizing energy consumption, cost, and greenhouse gas emission while maintaining thermal comfort and indoor air quality. The problem is formulated as a constrained convex optimization statement. Specifically, the thesis proposes three optimization-based control frameworks that are verified in simulation testbeds (with state-of-art simulation software and numerical models with MATLAB and Python). The three methods apply setpoint control on the room- and aggregate (building)- level devices and have achieved a 20% to 50% reduction in the peak load demand and greenhouse gas emission in simulation testbeds.

In addition to simulation, onsite experiments are conducted. One of the three simulation-based setpoint control frameworks is implemented in two MIT classrooms. Throughout the eight experiment sessions, a significant amount of commissioning of HVAC, software, and hardware is completed. This experimental verification has demonstrated a nearly 50% savings on greenhouse gas emissions and showcased the power of data-driven control methods in real-life settings.

Although we have witnessed the successes in both simulations and experiments, the results presented in the thesis are preliminary and only serve as a proof of concept. There are still plenty of areas worth further investigation to fully materialize and implement these methods.

Thesis Supervisor: Leslie K. Norford
Title: Professor of Building Technology

Thesis Reader: Steven B. Leeb
Title: Professor of Electrical Engineering and Computer Science and Mechanical Engineering

Acknowledgments

The past two years and a half have been fruitful, challenging, and self-fulfilling. With coursework in both computer science and building technology and two research projects related to building and energy sectors, I have explored a brand-new area where data-driven solutions are combined with traditional optimization and control frameworks and implemented in the building sector. In my steep learning path, I would like to acknowledge all my mentors and colleagues who have supported me along the way and witnessed my growing research capacity.

First and foremost, I would like to thank my advisor, Prof. Les Norford, for his continuous support and guidance. I appreciate the significant flexibility he offered for the research project selection and potential research directions so that I am encouraged to explore the areas of my interest. His patience and insightful feedback enable me to discuss my insights and express my concerns without hesitation.

I am also grateful for working with Dr. Kevin Kircher, the postdoc researcher on my team. When I was first introduced to the project with little domain knowledge, it was he who explained the concepts and showed me the approach step by step. Without him, I would not have picked up all the techniques and skills so rapidly and smoothly. I also appreciate Prof. Steven Leeb for offering me the opportunity to pursue a dual degree in EECS and Building Technology by introducing me to the project.

I would also like to thank members of my research group and my peers in the BT community. Thanks to Dr. Jeremy Gregory for offering the opportunity to work with the MIT campus on the energy-saving initiative. Thanks to Dr. Jasmina Burek and Dr. Subhro Das for setting up the EnergyPlus simulation model and establishing machine learning frameworks. Special thanks to Francis Selvaggio and other staff members from the Department of Facilities, as well as Steven Lanou from the Office of Sustainability for providing tremendous help in setting onsite experiments on campus. Thanks to Ray Gethers, Tyler Perron, and other engineers from Schneider Electric for consistent and prompt hardware and software support. Thanks to Stella Zhujing Zhang for walking me through and helping me set up the EPPY simulation. Thanks to Eduardo Gascon Alvarez for helping out with the Ladybug simulation. Thanks to Dr. Alpha Arsano for finding weather data sources for me.

Lastly, I would like to thank my parents and significant other as well as my friends who have always stood by my side, inspiring, supporting, and encouraging me to step out of my comfort zone and make a real difference at MIT.

Although the journey has come to an end, and I am ready to step into the next phase of my life, the challenges, enlightenment, and memorable experiences at MIT are the treasure that I will bring along the way.

Contents

1	Project Introduction	21
1.1	Background	21
1.2	Project Overview	21
1.2.1	Project 1: MIT Campus with Office of Sustainability	22
1.2.2	Project 2: A Nation-wide Utility Company	23
1.2.3	Team Contribution	24
1.3	Thesis Organization	25
2	Research Background	27
2.1	Demand Side Management	27
2.2	Thermostatically Controlled Appliances (TCAs)	28
2.3	Modeling Methods Overview	29
2.4	Objectives	30
2.5	Methodology Overview	31
2.5.1	MIT Classrooms @Building66	31
2.5.2	1R1C Thermal Capacitor and Resistor (RC) Model (Coupled and De-coupled)	32
2.5.3	Rule-based Control	34
2.5.4	Model Predictive Control (MPC)	34
2.5.5	Machine Learning Surrogate Models	35
2.5.6	“Predict, Perturb and Optimize” Scheme	35
2.5.7	Simulation Testbed	35
3	Three-step Peak Load Demand Reduction	37
3.1	Background	37
3.2	Methods	38
3.2.1	Baseline and Prediction	38
3.2.2	Load Perturbation	38
3.2.3	Load Shifting Optimization	39
3.3	Simulations	40

3.4	Results and Discussions	42
3.4.1	Learning a Perturbation Model with Linear and Non-linear Algorithms	42
3.4.2	Setpoint Optimization and Validation	45
3.5	Future Work	46
4	Two-step GHG Emission Minimization	49
4.1	Background	49
4.2	Methodology	50
4.2.1	EnergyPlus	51
4.2.2	Machine Learning Model Exploration	52
4.2.3	Convex Optimization	54
4.2.4	EnergyPlus Simulation Verification	56
4.3	Results and Discussion	56
4.3.1	Machine Learning Model Performance Comparison	56
4.3.2	Optimization Results	59
4.3.3	Load Minimization Verification	62
4.4	Future Work	64
5	Model Predictive Control for A Multi-objective Optimization	67
5.1	Background	67
5.2	Methods	68
5.2.1	Thermal Dynamics: Resistance Capacitance (RC) Network Model	68
5.2.2	Indoor Air Quality(IAQ): CO ₂ Mass Balance Model	70
5.2.3	HVAC Power Model	70
5.2.4	Dedicated Outdoor Air System (DOAS) Power Model	71
5.2.5	Grid Carbon Intensity	71
5.2.6	Electricity Wholesales Price	72
5.2.7	Occupancy Schedule	72
5.2.8	Weather Forecast	73
5.2.9	Objective and Constraints	73
5.3	Convex Optimization with Non-convex Simulation	74
5.3.1	Convex Relaxation: Discretize Differential Equations	75
5.3.2	Convex Relaxation: Small Time Interval Approximation	75
5.3.3	Convex Relaxation: Linear Approximation	76
5.4	Results and Discussion	76
5.4.1	Convex vs. Non-Convex Simulation Comparison	76
5.4.2	Model Response	79
5.4.3	Greedy Optimization vs. Optimization with Foresight	80
5.4.4	Multi-objective Pareto Front	80

5.4.5	Final Design Solution	83
5.4.6	Model Robustness	83
6	Experiment Setup and Design	85
6.1	Test Room Configuration	85
6.1.1	HVAC Equipment Setup and System Configuration	85
6.2	Experiment Design and Methodologies	86
6.2.1	Control and Baseline Test Sequence Design	86
6.2.2	Experiment Protocol	88
6.2.3	Data Requirement	89
6.2.4	Experiment One: Twin Test Room Validation	90
6.2.5	Experiment Two: Optimal Setpoint Implementation	91
6.2.6	Experiment Constraints	92
6.2.7	Two Optimization Framework and Targets	93
7	HVAC System Commissioning	97
7.1	Inability to Maintain the Setpoints	97
7.2	Constant Supply Air Volume (CAV) Operations in Variable Air Volume (VAV) Systems	98
7.3	Limited Cooling Capacity of Chilled Beams	99
7.4	Untraditional Control Sequence for Combined Systems	100
7.5	Varying Supply Air Setpoint of AHU	101
7.6	Frequent Oscillation of Supply Air Temperature	103
7.7	Sensor Placement and Calibration Susceptible to Human and Instrumental Errors	103
7.8	Inconsistent Observed Thermal Response	105
7.9	Unrobust Data Pipeline	106
8	Experiment Results and Discussion	107
8.1	Baseline, Optimal Control, and Rule-based (RB) Benchmark Setup	107
8.1.1	Load Calculation Example	109
8.2	Comparison of Baseline Estimation Methods	111
8.2.1	Validation 1: Parallel Test	111
8.2.2	Validation 2: Sequential Test	116
8.2.3	Validation 3: Scaled Predicted Baseline	119
8.2.4	Validation 4: Machine Learning Driven Baseline	121
8.3	Optimized Setpoint Control Results	121
8.3.1	Load Comparison	122

8.3.2	Relationship Between Room Loads and Outdoor and Indoor Temperatures	124
8.4	Future Work	126
9	Methods Comparison and Scalability and Implementation Discussion	127
9.1	Methods Comparison	127
9.1.1	Performance	127
9.1.2	Modeling Complexity and Scalability	128
9.1.3	Data Requirement	129
9.2	Machine Learning Performance Sensitivity Analysis on Data Scale	130
9.2.1	Data	130
9.2.2	Methods	131
9.2.3	Results	134
9.3	Campus Building Operations	136
9.3.1	Demand Side	136
9.3.2	Supply Side	139
10	Conclusion and Future Work	143
10.1	Conclusion	143
10.2	Future Work	144
A	Experiment Protocol and Setup	147
A.1	Typical Weekend Test Setup ¹	147
A.2	Illuminance and Temperature Sensors (HOBO) Placement	149
B	Experiment Detailed Documentation and Results	151
B.1	Experiment Documentation	151
B.1.1	Sep. 8 (Cooling Season)	151
B.1.2	Oct. 10 – Oct. 11 (Cooling Season)	152
B.1.3	Oct. 16 – Oct. 17 (Cooling Season)	152
B.1.4	Oct. 23 – Oct. 24 (Shoulder/Heating Season)	153
B.1.5	Nov. 20 – Nov. 21 (Heating Season)	153
B.1.6	Nov. 25 – Nov. 27 (Heating Season)	154
B.1.7	Dec. 26 – Dec. 30 (Heating Season)	155
B.1.8	Jan. 1 – Jan. 3, 2022 (Heating Season)	156
B.2	Room Load Sample Calculation (Chapter 8)	156
B.2.1	FTR and Chilled Beams	156
B.2.2	VAV	157

¹Depending on the experiment date, protocol may vary.

B.3	Predictive Baseline Load Calculation and Load Saving Estimation (Chapter 8)	158
C	Documentation of External Sources	159
C.1	Psychrometric Optimization for Supply Air Dry Bulb Setpoint for Buildings 18, E25, and 66	159
C.2	Sequence of Operations for HVAC Equipment (from DoF and Schneider Elec- tric)	160
C.3	Sensor Specifications and Documentations	163
C.3.1	Flowmeter Sensor: U1000MKII-FM Fixed Clamp-on, Ultrasonic Flow Meter	163
C.3.2	Temperature Sensors: Platinum Resistance Thermometers (PRT) . . .	163
C.4	EcoStuxture Setup and Site Navigation	164
C.5	Documentation of BMS and Other Databases Availability for Campus Buildings	169
D	Source Code	173
E	Documentation of EnergyPlus Simulation	175
E.1	Room 154 Model Parameters	175

List of Figures

1-1	Average hourly electricity load during typical day by region, selected months .	24
2-1	Illustration of peak clipping (curtailment) and load shifting	27
2-2	Air-conditioners operate on thermostatic control	29
2-3	MIT campus map and Building 66	33
2-4	1R1C network model illustration	33
3-1	A basic unit in simulation test bed, consisting of five zones (not drawn to scale)	41
3-2	Individual unit HVAC load for 10 simulated zones for a selected one-week period, with outdoor temperature	42
3-3	Room temperature dynamics with setpoint perturbation schedule for a selected zone	42
3-4	Model parameter tuning for temperature threshold (a) and memory (b) on training and validation sets	43
3-5	Predicted and actual perturbed loads for the validation dataset	44
3-6	Parameter tuning for model memory on the validation set with the neural network model	45
3-7	Predicted and actual perturbed load on the validation set from the neural network model	45
3-8	Validated and optimized perturbed loads with their setpoint perturbation for a selected week	46
3-9	Baseline, optimized, and validated aggregate loads for a selected week (a), and a selected day (b) using the simulation testbed	47
4-1	Optimized 24-hour temperature setpoint prediction system	50
4-2	An illustration of a 3-layer MLP model	55
4-3	Model performance comparison on a 24-hour forecasting horizon from the test set	57
4-4	Actual and predicted (with ARX model) hourly heating load trend in the heating season	57

4-5	Hourly heating energy associated with heating setpoints (a) and outdoor temperature (b)	58
4-6	(a) Optimized setpoints vs. grid MOER for test case 1; (b) Optimized setpoints vs. room occupancy status for test case 1	60
4-7	(a) Optimized setpoints vs. grid MOER for test case 2; (b) Optimized setpoints vs. room occupancy status for test case 2	61
4-8	(a) Optimized setpoints vs. grid MOER for test case 3; (b) Optimized setpoints vs. room occupancy status for test case 3	62
4-9	Validated vs. baseline GHG emissions. (a): test case 1; (b): test case 2; (c): test case 3	63
5-1	Convex vs. non-convex simulation result comparisons	77
5-2	MPC response to variation in grid carbon emission intensity	78
5-3	Control decision with different optimization horizons	79
5-4	Pareto front of multi-objective optimization	82
5-5	Model robustness	82
6-1	Building 66 floor plan	86
6-2	Room 154 BMS control dashboard	87
6-3	Classroom configuration	90
6-4	[Room 154] Cooling, heating setpoints and actual room temperature for 24-hour horizon on Sep. 23	93
6-5	[Room 154] The valve positions of VAV reheat coil and chilled beam on Sep. 23	94
6-6	[Room 154] Heating and cooling setpoints on Oct. 24 in the heating season	95
6-7	Hypothetical dual setpoint control with a varying setpoint deadband, an example for heating season	95
7-1	[Room 160] Setpoints and room temperature for 24-hour horizon on Sep. 8	98
7-2	[Room 160] Room and chilled beam chill water valve position on Sep. 8	100
7-3	[Room 160] Setpoints, room temperature (top), VAV supply air temperature and supply air flowrates (bottom) on Nov. 9	102
7-4	[Room 160] VAV supply air temperature and its reheat valve position on Sep. 11	103
7-5	[Room 154] FTR valve position and hot water flowrate comparison	104
7-6	VAV loads of Rooms 154, 156, and 160, during the baseline experiment weekend of Oct. 9 - Oct. 10	106
8-1	[Room 154] Room temperature, heating setpoints, and cooling setpoints for baseline (Nov. 25) and optimal control (Nov. 26) days	108

8-2	[Room 154] Room temperature, heating setpoints, and cooling setpoints for RB benchmark (Nov. 29)	109
8-3	[Room 154] An example of effective VAV airflow for heating and cooling . . .	110
8-4	[Room 154] Heating and cooling load profiles from FTR, chilled beam, and VAV, and outdoor temperature for baseline (Nov.25)	110
8-5	Baseline cooling load (only VAV) comparison in Rooms 154 and 160 on Aug. 18 and Aug. 19 (a) in the cooling season, and on Oct. 16 and Oct. 17 (b) in a shoulder season	112
8-6	Inside and outside conditions of Rooms 154 and 160	113
8-7	Light intensity comparison of Room 154 and Room 160 measured from the middle glass panel on Nov. 25 - Nov. 27	114
8-8	Measured hot water flowrates through FTR in Rooms 154 (a) and 160 (b) on Nov. 25 - Nov. 27 (heating season)	115
8-9	[Room 154] (Same as Figure 8-1) Room temperature, heating setpoints, and cooling setpoints for baseline (Nov. 25) and optimal control (Nov. 26) days .	117
8-10	Heating and cooling load profiles of FTR, chilled beam, and VAV, and outdoor temperature for Room 154 baseline on Nov. 25 (a) (Same as Figure 8-4) and optimal control on Nov. 26 (b)	118
8-11	Correlation of HVAC reheat load (a) and FTR load (b) between Room 160 and Room 154	120
8-12	Comparisons of the average electricity carbon intensity for baseline and optimal control tests	121
8-13	[Room 154] VAV reheat and FTR load comparison among predicted baseline, optimal control, and RB benchmark	122
8-14	Distribution of temperature deviations of two baseline, optimal control, and RB-benchmark experiments	123
8-15	[Room 154] Average room temperature in baseline tests and optimal control tests	125
8-16	[Room 160] Outdoor temperature vs. heating loads	125
9-1	Ranges of hourly cooling loads for the 38 buildings on campus after removing buildings with outlier data	131
9-2	Cooling load time series for the 38 buildings, plotting before (blue) and after the Hampel correction (orange)	132
9-3	Distribution of points corrected by the Hampel filter across all buildings, by month of year and hour of day.	133
9-4	Hourly cooling load prediction of Building 66 with AR (6, 1) model for heating (a) and cooling (b) seasons	135

A-1	Signage examples	148
A-2	HOBO (illuminance and temperature) sensors placement location in Room 154 (a) and Room 160 (b)	150
C-1	Psychrometric optimization for supply air dry bulb setpoint for Buildings 18, E25, and 66	159
C-2	Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160	160
C-3	Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160 (continued)	161
C-4	Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160 (continued)	162
C-5	Flowmeter Specifications	163
C-6	Global Protect VPN and Gateway	164
C-7	Workstation login type	165
C-8	Main panel	165
C-9	Site Navigation	166
C-10	Floor plan overview, AI project dashboard	166
C-11	Setpoint adjustment dashboard and BTU meter reports	167
C-12	BTU data download	167
C-13	Requested data summary	168
C-14	Forcing setpoints to allow a varying setpoint deadband	168
C-15	BMS and other databases availability for campus buildings	170
C-16	BMS and other databases availability for campus buildings (continued)	171
C-17	BMS and other databases availability for campus buildings (continued)	172
C-18	BMS and other databases availability for campus buildings (continued)	172
E-1	[Room 154] EnergyPlus Simulation model	175

List of Tables

2.1	Summary of simulations and experiments	32
3.1	Model parameters for the simulation testbed	41
4.1	Model performance comparison	57
4.2	Baseline and validated optimal GHG emission comparison	64
5.1	Model variables, parameters, constraints, and objectives	69
6.1	Baseline and optimal control test schedule for a two-day test period (a) and a three-day test period (b)	89
6.2	Data types and databases for experiment setup and result analysis	90
7.1	Sensors, meter types and their usages in onsite experiments	104
8.1	Evaluation of thermal comfort for baseline, optimal control, and RB-benchmark test days	124
9.1	Model performance (RMSE) on cooling hourly load prediction of Building 66	134
9.2	Summary of AR(6) model sensitivity to different prediction scales	136
9.3	Five factors for evaluating the eligibility of space-/building-level participation in the data-driven setpoint control project	137
B.1	Sample data	156

Chapter 1

Project Introduction

1.1 Background

Buildings and building construction sectors are together responsible for nearly 36% of global energy consumption, 70% of electricity consumption, and 40% of CO₂ emissions [1, 3]. It is predicted that building energy usage will continue to grow globally over the next decade, associated with a higher standard of living enabled by a wider access to household electric appliances, such as Heating, Ventilation, and Air-Conditioning (HVAC) systems that achieve improved occupant thermal comfort and indoor environmental quality [82].

In the meantime, the Biden administration has set a target of a 50% reduction in 2005 GHG emissions by 2030 and net zero emissions by 2050 [81]. Achieving these targets will require deep decarbonization across all sectors. This is particularly challenging in the building sector given that the existing building stock is large and there is relatively low turnover. Thus, to achieve such deep decarbonization targets the US must simultaneously create energy-efficient new buildings while lowering energy consumption in existing buildings.

Building scientists have actively worked on saving energy usage by developing methods that improve building performance and energy efficiency, such as installing better thermal insulation [83], developing advanced control algorithms for operating lighting and HVAC systems [61], and performing life-cycle analysis for building construction [19].

1.2 Project Overview

The thesis includes two main projects from the perspectives of two stakeholders: (1) energy end user, and (2) energy provider. While the goals of two projects vary, both project sponsors would like to regulate and optimize energy usage behaviors to maximize their economic and environmental outcomes. Therefore, the methods explored and deployed share a certain level of similarity in nature.

1.2.1 Project 1: MIT Campus with Office of Sustainability

The first project is with MIT. In response to the clean energy initiative, MIT, as a leading academic institution collaborating with industry pioneers, has invested capital to renovate campus buildings and upgrade heating and cooling system control to improve building energy efficiency. This project, sponsored by the MIT Office of Sustainability, aims at further investigating efficient operations of heating, cooling, and air conditioning (HVAC) systems on the MIT campus through thermostatic control, minimizing greenhouse gas (GHG) emissions from gas and electricity consumption while maintaining satisfactory thermal comfort to students, faculty members, and staff.

The campus environment offers unique opportunities for experimenting and implementing GHG emission savings, but also poses special challenges:

On the one hand, the MIT campus has its independent energy management task force, led by Department of Facilities (DoF), Office of Sustainability, and MIT Energy Initiative, consistently monitoring campus energy usage. Because MIT has autonomy in regulating its energy consumption, implementation of adjustments would face few obstacles and can be integrated into MIT's long-term energy emission reduction blueprint. In the meantime, unlike residential or commercial buildings, educational institutions follow a relatively predictable building operation schedule according to course schedules and lab activities. This is conducive to regulating energy usage as the load pattern is more predictable and controllable.

On the other hand, MIT both purchases electricity from the market (Eversource) and generates electricity through a tri-generation plant. The DoF sources energy (i.e. when to use onsite or offsite electricity) based on market information (electricity price) and campus energy demand, which involves a complicated modeling and optimization process. Therefore, an effective GHG emission reduction solution should carefully consider its compatibility with existing campus energy supply and procurement strategies. In addition, different space types on campus require tailored design of control methods to fulfill their specific indoor environmental requirements. For example, labs tend to have more rigorous environmental controls than offices and classrooms to ensure occupants' health and safety. Before scaling up solutions to the entire campus, a comprehensive method evaluation on different space types should be conducted to minimize adverse effects on space users.

MIT is not the first that initiates a campus-wide building HVAC operation and thermal comfort project. Back in 2017, UC Davis launched TherMOOstat website¹) for their campus thermal comfort initiative, where students can input their feedback on room temperature in most campus buildings [9]. Similarly, this MIT project also includes a thermal comfort survey that allows students to provide real-time thermal comfort feedback. The survey results (i.e. preferred room temperature range) could also inform room setpoint design as a critical element of the GHG minimization project.

¹<https://thermoostat.ucdavis.edu/>

1.2.2 Project 2: A Nation-wide Utility Company

The second project is with a large electricity and natural gas utility supplier, in collaboration with the Department of Electrical Engineering and Computer Science (EECS) at MIT. The project's ultimate goal is to develop a self-driven load control model that can autonomously smooth load demand for an aggregate (large-scale community).

From International Energy Agency (IEA)'s Smart Grids Technology Roadmap in 2011, an increasing peak load and rising electricity consumption are the two major trends that impact future grid operations [46]. According to EIA's analysis on the average hourly electricity overview in Figure 1-1, electricity consumption cycles every day with a trough occurring around 5 am and the one or more peaks occurring during the daytime, before dropping during the late evening hours. While the daily profile varies seasonally and geographically, the daily demand variation can be very significant, as illustrated in Figure 1-1, up to 50%. To fulfill grid peak load demand, utilities ramp up stand-by small-capacity power plants that run on less energy-efficient and environmentally-friendly fuel such as diesel and coal. It is also economically unattractive to maintain operations of a stand-by fleet because they only operate during peak load hours [76, 26]. Hence, an increase in the peak load demand will further impose additional stress on power generation systems designed with an output capacity limit, causing a reduction in power generation efficiency, the rise of grid carbon intensity, emissions of toxic air pollutants such as NO_x and SO_2 , and an increase in electricity transmission loss.

To deal with hundreds and thousands of grid end users from diverse sectors, the utility supplier would like to explore less user-intrusive, cost efficient, and easy-to-implement solutions, monitoring and regulating users' energy demand profile. While Internet of Things (IoT) enabled technologies such as Nest Thermostats developed by Google learn optimal device control decisions by collecting and analyzing user behavior information, it is financially and computationally expensive to deploy the hardware and establish a giant control hub to provide data-driven solutions for all end users. Furthermore, grid users may be sensitive to releasing their personal information such as daily schedule to the external network, so a non-intrusive and locally distributed method is more suitable.

This project aims to manage a wide range of loads for various end users, such as HVAC systems in commercial buildings, electric vehicles, refrigerators and water heaters in residential buildings, and water pumps. Collaborating with MIT EECS researchers, this project is composed of two main aspects: first, to develop a non-intrusive load monitoring algorithm that can detect and filter out load patterns generated from different usage of appliances; second, to construct a distributed control method that regulates individual power usage. The researchers in the EE department are in charge of the first part, and the Building Technology group is mainly involved in the second half.

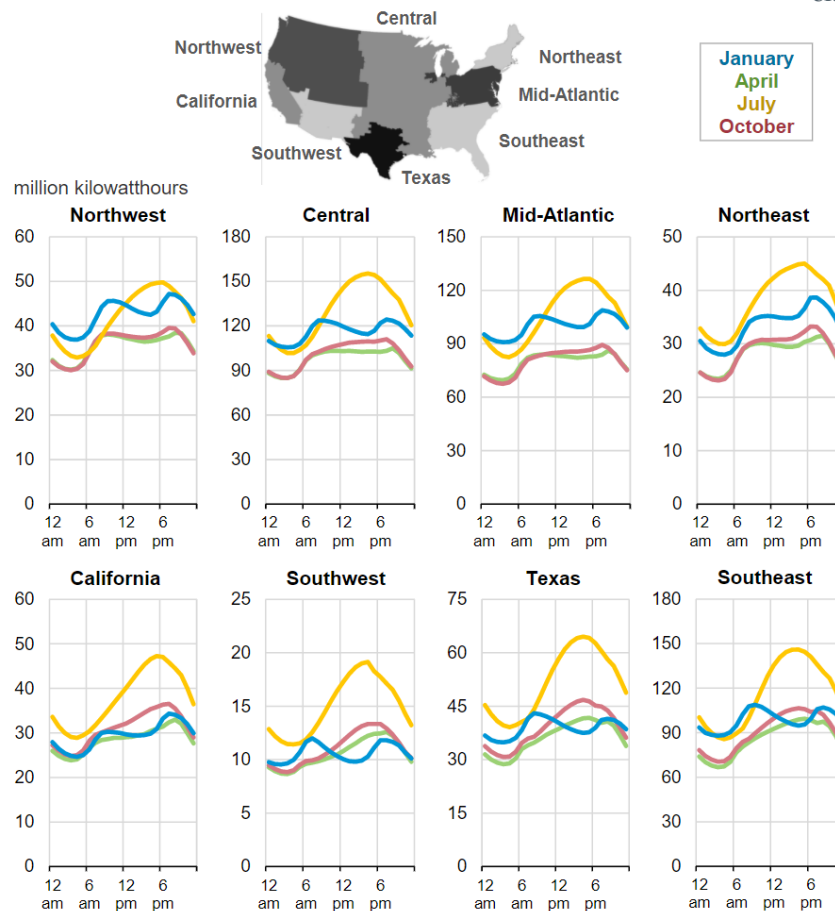


Figure 1-1: Average hourly electricity load during typical day by region, selected months [6]

1.2.3 Team Contribution

The methodologies and results presented in the following chapters are a result of team efforts. In the first MIT project, the research group includes Dr. Jeremy Gregory (Executive Director of the Climate and Sustainability Consortium), Prof. Les Norford (Professor in Building Technology Group), Dr. Kevin Kircher, (postdoc researcher in EECS Department), Dr. Jasmina Burek (former postdoc researcher at MIT Materials Systems Laboratory and currently Assistant Professor at University of Massachusetts Lowell), Julia Wang (M.Eng Student in EECS Department), Dr. Subhro Das (Research Staff Member at the MIT-IBM Watson AI Lab). In addition, MIT DoF and Office of Sustainability, and engineers from Schneider Electric offered un-negligible support to facilitate methodology development and hardware deployment in the experiment phase. The second project involves members from MIT EECS Department – Dr. Kevin Kircher, Adedayo Aderibole (Ph.D. student) Prof. Steven Leeb, and Building Technology Group from Architecture Department – Prof. Les Norford.

1.3 Thesis Organization

The thesis chapters are structured in four major sections:

- **Section 1** - Chapters 1 and 2: define research problems and present the project overview with literature review in related fields.
- **Section 2** - Chapters 3 to 5: propose three methods to address the research statements and evaluate the results in simulation environments.
- **Section 3** - Chapters 6 to 8: document the procedure for onsite experiments, identify and propose solutions for observed hardware and software issues, assess the effectiveness of the selected approach in an experimental setting.
- **Section 4** - Chapters 9 and 10: compare and contrast the three simulation- based approaches and their practicality and feasibility for large-scale implementation, address other considerations, and summarizes the thesis with suggestions of future work.

Chapter 2

Research Background

2.1 Demand Side Management

Trends toward increasing energy demand, especially electricity, and its rising peak load, observed over the past several decades and forecast for the near future, has inspired scientists, economists, politicians and other stakeholders to design a set of mechanisms that can regulate energy demand. Demand side management (DSM) is one technique. As its name suggests, DSM provides a list of strategies that utility companies, power suppliers, and transmitters can utilize to control power consumption profile of end users in response to energy supply conditions [64]. DSM has two common goals to achieve [52]: (1) peak load curtailment, and (2) load shifting, as illustrated in Figure 2-1. Load curtailment shaves the load during the anticipated peak-load periods. Load shifting brings power consumption during peak-load periods to off-peak load periods. With load curtailment and load shifting, both energy supply and demand sides would enjoy significant savings as the electric grid becomes more efficient and sustainable [60]. To be more specific, the US Energy Information Administration (EIA) estimates that eliminating demand peaks could lead to a huge reduction in transmission and distribution losses, which is equivalent to 5% of electricity transmitted and distributed in the US in 2015 through 2019 [5].




Figure 2-1: Illustration of peak clipping (curtailment) and load shifting [37]

To achieve the DSM goals, two major types of methods are involved: dynamic pricing

and direct load control (DLC). The basis of dynamic pricing is that it considers the rising cost of electricity generation during peak hours and, therefore, adding surcharges to bills based on consumers' usage [21]. By setting a higher price at peak demand and a lower at trough, it incentivizes consumers to voluntarily shift energy consumption activities to an off-peak period [20, 2, 7]. However, establishing dynamic pricing and a reimbursement program requires active and voluntary customer involvement [17], which may not be as effective as expected because end users may not 1) be aware of such a price difference mechanism, and 2) or care about a minor increase in monthly recurring bills. On the other hand, DLC allows utilities, or an aggregator, to remotely control the operations and energy consumption of customers' certain appliances, such as lighting, HVAC equipment, refrigerators, and pumps [85, 77]. For example, the state of California, collaborating with PG&E, initiated the Energy Incentive Program to reimburse customers who choose to adjust their energy usage, i.e. by turning air-conditioners off during peak load hours for an agreed amount of time [8]. To accurately regulate end-user loads while not impacting consumers' comfort, it is important to deploy advanced automation, control, and monitoring technologies that make DLC less hindering to the customers [77]. The two projects presented in the thesis utilize DLC to achieve peak load reduction and GHG emission savings¹.

2.2 Thermostatically Controlled Appliances (TCAs)

In residential and commercial units, DLC targets background electrical loads, such as heating, ventilating, and air-conditioning units (HVACs), refrigerators, and electric water heaters. This is because these loads operating in the background are relatively flexible and can be easily controlled and adjusted for load control purposes, without sacrificing too much user experience. For example, the setpoint temperature of an air-conditioning system can be modified to adjust power consumption without introducing excessive occupant thermal discomfort. Due to the nature of background electrical loads, most appliances are modeled as thermostatically controlled appliances (TCAs). TCAs refer to on/off electrical devices that maintain their operating temperatures within a range of a user-defined setpoint. Various studies have investigated using different TCAs to provide DLC [71, 25, 65]. For both projects in this thesis, the research focus is on HVAC systems, which are a major power consumption component in many power systems. In its annual report, EIA estimates that in 2020, electricity used for space cooling by the U.S. residential and commercial sectors is about 10% of total electricity consumption [12].

To illustrate the operation of a thermostatically controlled air-conditioner, Figure 2-2 shows its on/off schedule to maintain its indoor temperature near the setpoint T_{set} , with a

¹As noted in Chapter 1, MIT partially manages its energy generation and supply, so the campus can be considered as a microgrid where MIT DoF is the "utility" regulating energy supply and demand

deadband of 2ϵ . In this figure, the air conditioner is initially off so the indoor temperature drifts upward until it reaches the maximum temperature allowed by the thermostat, $T_{set} + \epsilon$. The thermostat then switches the air conditioner on, driving the temperature downward. When the temperature reaches the minimum allowed temperature, $T_{set} - \epsilon$, the thermostat switches off the air conditioner and the process repeats itself.




Figure 2-2: Air-conditioners operate on thermostatic control

To operate TCAs in the desired manner, researchers designed control algorithms based on their thermostatic properties. Common control algorithms include rule-based control, PID control, model predictive control, model-free control, and reinforcement learning [62, 15, 78, 56]. Some control methods also integrate other mechanisms, for example priority stack control that assigns an urgency of fulfilling the load demand [43, 74, 73]. Sanandaji et al. [74] also added anti-short-cycling mechanisms to prevent deterioration of equipment.

Numerous studies have investigated DLC using different methods to address peak load shifting at different scales. While the concept of DLC has existed for decades, most studies develop new methods and verify them in a simulation environment. Only a few conducted field experiments, though with simpler control policies [35], and in a small and manageable site [75] because onsite experiments incur significant cost and could interrupt normal building operations. The difficulty in doing onsite experiments have two implications. First, an accurate thermal simulation model is critical to understanding the thermal properties of the modeled device/space. Second, an experimentally verified method is extremely valuable to the entire building research community. In this thesis, we develop control algorithms for HVAC systems to provide DLC in simulation settings and verify their effectiveness in simulations and experiments. The following section introduces the modeling methods explored. A good simulation model can not only provide insights into building physics, but also establish a platform to exercise building control algorithms.

2.3 Modeling Methods Overview

A variety of thermal modeling methods have been investigated and applied to TCAs [48]: (1) an analytical and physical model with numerous building parameters and complicated differential and algebraic equations, replicating the exact physical heat transfer process

defined by the modeling boundary; (2) the Resistance Capacitance (RC) networks or the semi-physical thermal circuit model, analogous to an electrical circuit to monitor heat flow as electric current; and (3) a data-driven machine learning model that learns thermal behaviors based on knowledge extracted from past observations. The three modeling methods have their own application environments and limitations. For example, the analytical method enabled by simulation software, such as EnergyPlus, gives an accurate representation of the thermal performance of a modeled space but is not easily scalable to a large community, as the expense of scaling up models drastically increases with explosion of modeling parameters [31]. The RC network approach simplifies the heat transfer process by spreading the effect of the heat transfer to several key modeling parameters termed as thermal capacitors and resistors [48]. The few parameters and simplified modeled heat transfer processes ensure an easy and efficient scale-up. On the other hand, a machine learning model is typically trained on data obtained from simulation results or in real experiments, but the downside is that in practice the availability and quality of data are not guaranteed.

Studies with various modeling and control methods also explored the impact of modeling scales on their ability to reduce peak load demand. Lee et al. [50] used a stochastic queuing model to control the operations of air conditioners in a 449-unit apartment complex, which achieved 24% peak load reduction. Nghiem et al. [59] applied decoupled first-order linear systems (1R1C Resistance-Capacitance Model) and proposed a two-level control architecture on 5 and 20 loads, which reduced the peak aggregate demand by up to 40%. Various reinforcement learning techniques, such as Q-learning, BRL and W-learning, have been recently investigated to reduce peak demand with electric vehicles [57] and electric water heaters [16].

Notwithstanding the considerable amount of research studies in peak power demand reduction, a systematic discussion of each approach and their application environments has not been addressed. First, when we review new simulation-based control algorithms, it is unusual to find the results with new approaches benchmarked with the existing ones or certain classical control models, or the quantification of the benefits of implementing the new model. Second, most of the new algorithms are developed in simulation testbeds, assuming perfect knowledge parameters of the model and control of model inputs, such as building operation schedule, room thermal properties, and occupant behaviors. The lack of onsite experiments fails to validate performance in real-life, which can be very different from simulation-based results.

2.4 Objectives

In proposing and testing new modeling and control methods to address the two project goals, this thesis also includes:

- Evaluation of different modeling and control approaches in various application environments in terms of their modeling efficiency, scalability, robustness, and their capacity of reducing peak power consumption or GHG emission, using simulation testbeds.
- Assessment of results from experiments conducted in MIT classrooms and their implications.
- Overview of hardware and software requirements and experiment protocol for a successful implementation of field experiments on the MIT campus.
- Discussion of method scalability to large-scale instrumentation, and implications and suggestions for future experiment implementation.

Table 2.1 summarizes simulation methods and control algorithms involved in each proposed study, followed by a brief discussion of methodologies and experiment setup. The purposes of the studies summarized in Table 2.1 are either (1) to reduce total GHG emissions or (2) to minimize peak load demand, associated with HVAC system operations through thermostatic control, while not sacrificing occupant thermal comfort and/or indoor air quality. The studies can be classified into three categories: (1) a benchmark model that uses a simple rule-based control algorithm, (2) a room-level simulation-based model, followed by experiments in MIT classrooms for optimal setpoint control, and (3) an aggregate-level simulation-based model.

2.5 Methodology Overview

2.5.1 MIT Classrooms @Building66

Building 66 is the major site for experiment instrumentation and HVAC system investigation (See Figure 2-3 map below). It is home to the Department of Chemical Engineering, with majority space used as research labs, offices, and classrooms. In this study, we chose two classrooms Room 154 and Room 160 for pilot research. We installed occupancy sensors and BTU meters to keep track of room occupancy status and thermal power consumption. In addition, building and room level data such as room temperature, supply air temperature, setpoints, and supply air flowrate are collected and maintained at Clockworks². The control platform EcoStruxture provided by Schneider Electric³ allows us to adjust room temperature setpoints for rooms on an hourly basis. In addition to the two-classroom experimentation, further analysis of other HVAC system types was conducted to address the feasibility and applicability of the solutions to campus-level instrumentation.

²<https://clockworksanalytics.com/>

³<https://ecostruxureit.com/> MIT has its own EcoStruxture workstation installed to access the the building management system

Table 2.1: Summary of simulations and experiments

	Scale	Simulation or Experiment	Control Algorithms	Purpose
Benchmark Model for MIT Classrooms @Building 66	2 classrooms	Onsite experiment	Rule-based control	Reduce total GHG emission from HVAC operations and maintain occupant thermal comfort
Room-level Model	1 classroom	Numerical simulation testbed (decoupled 1R1C model)	Model predictive control	Reduce total GHG emission and/or energy consumption, and maintain occupant thermal comfort and indoor air quality
MIT Classrooms @Building 66	2 classrooms	EnergyPlus simulation with machine learning surrogate models	"Predict and optimize" 2-step framework	Reduce total GHG emission from HVAC operations and maintain occupant thermal comfort
MIT Classrooms @Building 66	2 classrooms	Onsite experiment using simulation results	/	Reduce total GHG emission from HVAC operations and maintain occupant thermal comfort
Aggregate-level Model	Multiple buildings	Numerical simulation testbed (Coupled 1R1C) with machine learning models	"Predict, perturb and optimize" 3-step framework	Reduce peak load demand and maintain occupant thermal comfort

Such sensor and meter installations for experimentation rarely occur because they are time and labor consuming and involve various stakeholders across campus, such as the Department of Facilities and Schneider Electric, which initially developed the control platform. The customized control dashboards allow us to control the thermostat operations of dozens of spaces simultaneously⁴. This provides us with a unique research opportunity to validate our proposed method using real buildings, and to gain insights into the values brought by the proposed methods in real-world implementations.

2.5.2 1R1C Thermal Capacitor and Resistor (RC) Model (Coupled and Decoupled)

An RC network model is analogous to an electric circuit model as a numerical simulation method [27], as shown in Figure 2-4. This 1R1C model includes one capacitor, equivalent

⁴In Building 66, we only have room-level control, but in the other experiment site, NW23, a multi-zone level control dashboard has been provided to allow simultaneous control of multiple zones in the building.




Figure 2-3: MIT campus map and Building 66

to thermal capacitance of mass and air, and one resistor that lumps all heat transfer effect such as conduction and convection into one process. Studies have shown that a first-order model, albeit less accurate than a second-order model, is often sufficient to describe the heat transfer process within a space [51]; a higher-order model only brings diminishing marginal modeling accuracy while incurring larger modeling costs.




Figure 2-4: 1R1C network model illustration

In Figure 2-4, T_a represents ambient temperature; T is indoor air temperature assumed the same as mass temperature; thermal power \dot{Q}_c is heat injected by controlled thermal equipment, such as HVAC equipment; and thermal power \dot{Q}_e is heat injected by exogenous sources that contribute to thermal disturbance, which may include solar radiation, heat brought in by outdoor air, and occupant and lighting loads. Thermal capacitance and resistance (R and C) are properties of room construction materials. Using Kirchhoff's current law, we get

$$\dot{Q}_c + \dot{Q}_e = C \frac{dT}{dt} + \frac{T - T_a}{R} \quad (2.1)$$

The solution of this first-order model gives a linear representation of the system at time step $(k + 1)$

$$T(k + 1) = aT(k) + (1 - a)[T_a(k) + R(\dot{Q}_c(k) + \dot{Q}_e(k))] \quad (2.2)$$

where

$$a = e^{-\frac{\Delta t}{RC}}$$

Δt is a discrete time interval; \dot{Q}_c is positive in the heating season and negative in the cooling season.

When modeling multiple rooms within a building, we can treat the system either as coupled or decoupled. A coupled system indicates that heat transfer happens between rooms, for example, through heat conduction via connecting walls and convection through airflow movement between rooms. Thus, a coupled system requires an additional resistor between each room. A decoupled system, on the other hand, assumes each room transfers heat only with outdoor space, and is adiabatic to its adjacent rooms. This thesis includes two RC models, with one applying to room-level simulation and the other to aggregate building-level. Only a building-level model uses a coupled configuration to consider a space's heat transfer with its surroundings within a building.

2.5.3 Rule-based Control

The rule-based control algorithm is not the major focus of this project - it only serves as a benchmark for proposed method comparison. A rule-based control algorithm adjusts temperature setpoints according to occupancy status. The controller will be notified of the room occupancy schedule before deciding room setpoints. An unoccupied period permits a more flexible setpoint range, while an occupied period uses a Business-As-Usual setpoint profile (following the existing schedule implemented in the campus buildings.)

2.5.4 Model Predictive Control (MPC)

Traditionally implemented in other engineering fields, model predictive control is used to control a process while satisfying a set of constraints. More recently, it has been applied in building control systems due to its capability of dealing with dynamical factors in the heat transfer process and power grid behaviors [13, 55]. MPC is able to make decisions based on predicted future scenarios, and therefore is capable of planning an optimized control schedule for the near future. MPC shows strength in building HVAC operation control because it can optimize building energy consumption, cost, and greenhouse gas emission, based on the predicted conditions in the electricity market, energy usage, and space temperature dynamics. For example, MPC can leverage opportunities such as shifts in energy price and changes in grid carbon emission intensity to strategically manipulate the HVAC operation's schedule, which leads to minimum operating cost and GHG emission.

2.5.5 Machine Learning Surrogate Models

Machine learning models with different scale and prediction objectives are widely deployed to capture embedded information from the observed or simulated data points, such as power consumption and room occupancy count, and to make accurate predictions for building operations planning. Due to the nature of the prediction task, time series models are the primary candidates, framed into linear regression, autoregression with exogenous inputs, multi-layer perceptron (a simple form of neural network model), and long short-term memory (a type of recurrent neural network) structures [31, 11, 72]. The model inputs and outputs vary from each prediction task. The predictive power of an accurate surrogate model exempts the researchers from laboriously configuring complicated simulation models. Besides, it allows a seamless integration with the optimization framework that is usually less compatible with traditional simulation approaches. Furthermore, it generalizes and scales room-level or building-level findings to a larger application context, which fulfills the ultimate application setting for both projects.

2.5.6 “Predict, Perturb and Optimize” Scheme

This scheme is developed and formulated collaboratively with Dr. Kevin Kircher to optimize aggregate or room level peak load demand, energy costs, and GHG emission. It has two variants: (1) three-step control and (2) two-step control. Three-step is configured to provide aggregate-level setpoint control, including: (a) learning a baseline power prediction model, (b) identifying a model describing perturbations about baselines, and (c) embedding baseline predictions and perturbation model parameters in optimization that outputs optimal setpoint schedules. The two-step method is applied in room-level control and combines the step (a) and (b) into one process by directly predicting the resulting perturbed power consumption with the perturbed setpoints, disregarding the baseline reference.

2.5.7 Simulation Testbed

In addition to EnergyPlus simulation, we established numerical simulation testbeds in Python that consist of a system of 1R1C network models, integrated with one specific controller, such as MPC, or “Predict, Perturb and Optimize” approach, to model and control the thermal performance and response of individual spaces or the aggregate, by providing external conditions such as the ambient temperature and occupancy schedule. The simulation testbeds output hourly power consumption for each zone and whole building. Although testbeds would never perfectly represent the thermal performance of buildings in reality, they enable researchers to freely explore different temperature setpoint control strategies, not limited by constraints in reality regarding thermal comfort consideration, on/off equipment switch frequency concern and so on.

Chapter 3

Three-step Peak Load Demand Reduction

3.1 Background

This three-step approach is developed for reducing large-scale power grid peak load demand by adjusting setpoint temperatures of HVAC units. This approach can also be further extended to other devices controlled by setpoint temperatures, such as water heaters and refrigerators, or devices with energy storage capacity, such as electric vehicles [47].

Applying setpoint adjustment to manage load demand is not a new idea. Researchers have developed diverse methods to accomplish this goal, ranging from establishing model-based control for distributed optimization [49], to reinforcement learning for various and nonlinear loads [80, 22, 84].

While these methods demonstrate potential, they suffer from several major limitations. First, it is almost impossible to develop simulation models perfectly matching the thermal response of buildings, not to mention the stochastic occupant behaviors. The numerical models often have to rely on simplified assumptions that could generate performance gaps. Second, while reinforcement learning is model-free, it requires an extremely careful training and calibration process, and its computational complexity explodes when the algorithm is scaled up. In addition, implementation of the algorithm in real buildings requires considerable upfront setup time for the model to learn and explore building behaviors.

In comparison, the three-step approach is designed in such a way that it avoids the algorithmic and implementation challenges. It is easy to scale to aggregate load for various types of devices because it requires no device-level models, which also protects user privacy. It also needs minimal computation and sensing as the only information it gathers and processes is the aggregate power, which is independent of the number of devices involved. At the device level, it only requires the device to be connected to the network to receive messages for

setpoint adjustment.

3.2 Methods

The three-step approach includes the following: (1) establishing and learning a baseline prediction model, (2) identifying a linear model generalizing load perturbations about baselines, and (3) embedding baseline predictions and perturbation models in convex optimization to provide optimally control setpoint schedules.

3.2.1 Baseline and Prediction

The baseline aggregate power $\hat{P}_t(k)$ refers to the business-as-usual power consumption to track user-specified setpoint schedules for multiple units and devices at discrete time step k . A time-series forecasting model can be used to predict aggregate power under baseline operation. The model requires historical aggregate power measurement data and predictive features such as weather conditions (ambient temperature, relative humidity, and cloudiness) and temporal information (seasonality, week of day, and hour of day). Possible model structures include linear regression, linear time series, neural networks, and support vector machines. This chapter does not incorporate machine learning baseline models because we would like to test the validity of the approach first by comparing the load savings directly with the true simulated baseline without introducing confounding factors, such as accuracy of a prediction model. However, more time series load forecasting problems are explored in Chapter 4 and Chapter 9 where model comparison and data requirements are discussed in detail.

3.2.2 Load Perturbation

The perturbation is simultaneously applied to all user-specific setpoints, and the allowable perturbation is ± 1 °C. The perturbation could either increase or decrease the total aggregate power consumption $P_t(k)$. The difference between the baseline and perturbed aggregate power is the load perturbation $u(k)$, where

$$u(k) = P_t(k) - \hat{P}_t(k) \tag{3.1}$$

To generalize the load perturbation response, we implement a linear model which is compatible with our following optimization framework and find sufficient for the prediction task. Based on Taylor’s theorem, the load perturbations should respond approximately linearly to small setpoint (± 1 °C) adjustments, although the actual response is generally non-linear. To reflect the unique nature of cooling and heating loads, the linear model is structured differently from a common one where regression is performed to define model coefficients

between inputs and outputs directly without considering historical trend. Specifically, the power perturbation at time k depends on not only the current setpoint perturbation, but also the past history to account for the historical thermal dynamics. In the meantime, there exists the non-linear property of perturbed power due to the nature of the perturbation experiment - the perturbed power would become zero if the setpoint is lowered by 1 °C in heating season and increased by 1 °C in cooling season. To account for these factors, we formulated the problem such that the trend can be easily recognized by the algorithm:

$$u(k) = \alpha_1(k)\delta_T(k) + \dots + \alpha_m(k)\delta_T(k - m + 1) + \epsilon(k) \quad (3.2)$$

where $u(k)$, $\delta_T(k)$, and $\epsilon(k)$ are load and setpoint perturbations, and the error term at time k ; parameters $\alpha_i(k)$ vary with outdoor temperature:

$$\alpha_i(k) = \beta_i \max\{0, \theta(k) - \theta_{\min}\} \quad (3.3)$$

such that the model predicts $u(k) = 0$ whenever $\theta < \theta_{\min}$; and when above θ_{\min} , the predicted $u(k)$ increases linearly with $\theta(k)$.

Parameters β_1, \dots, β_m are fit by running a regression model; memory m and temperature threshold θ_{\min} are tuned by grid search, based on root mean square error (RMSE).

While a linear model is formally adopted in the framework, we also experimented with non-linear neural networks for time series analysis with similar predictive features except that the non-linear model does not involve very physical based specific structure. The linear network predicts perturbed loads for a 4-hour time horizon. The inputs are scaled to standard normal distribution before being sent to a multi-layer perceptron (neural network) with a hidden layer of size 80 units, followed by a ReLU non-linear function and a drop out rate of 0.5. The model is trained using mean absolute error (MAE) and evaluated with both MAE and RMSE as metrics.

3.2.3 Load Shifting Optimization

In the last phase, the baseline predictions and the perturbation model are embedded in the convex load-shifting optimization. The optimization framework can be extended to other objectives such as thermal comfort, energy costs, GHG emissions, and revenues from power system services such as frequency regulation. In this project, the optimization objective is to limit peak aggregate demand by regulating temperature setpoint perturbations $\delta_T(0), \dots, \delta_T(K - 1)$:

$$\begin{aligned} & \text{Minimize} && \max_{k \in \{1, \dots, K\}} P(k) \\ & \text{Subject to:} && |\delta T(k)| \leq 1 \end{aligned} \quad (3.4)$$

where $P(k) = P_t(k) + \epsilon(k) = \alpha_1(k)\delta_T(k) + \dots + \alpha_m\delta_T(0) + \epsilon(k) + \hat{P}_t(k)$.

This is a convex optimization problem with $K + 1$ variables and one constraint. It is reformulated as a linear program and solved by a standard optimization software, in this case, CVXPY [14, 29, 30].

3.3 Simulations

A simulation testbed for an aggregate community is developed to validate the proposed method. The model adopts a coupled 1 resistance 1 capacitance (1R1C) network model that accounts for heat transfer between the interior and ambient environment, as well as between the neighbors. The testbed enables users to define the number of basic units (Figure 3-1) to adjust the size of the aggregate. The basic simulation unit is composed of five zones in a pre-defined layout, with four exterior zones, and one interior zone. After the simulation size is decided, the algorithm randomly generates model parameters within a reasonable range, including R and C values, HVAC equipment coefficient of performance (COP), and equipment sizing factor, etc. Depending on the position of the zone, a window size and solar heat gain coefficient are assigned to each zone, and solar radiation generated from simulation software Ladybug [63] is also included as a zone feature. With the defined parameters, the testbed automatically sizes the HVAC heating and cooling capacity based on the Actual Meteorological Year (AMY) weather data. The simulation testbed is placed in Austin, TX and the entire simulation utilizes weather data in that area. In the meantime, to include not just the electricity load from HVAC systems, we incorporate the non-HVAC loads such as lighting, plug loads, and occupancy introduced loads by referring to the Cornell campus dataset provided in the ASHRAE Great Energy Predictor Competition [58]. All simulation parameters are summarized in Table 3.1. The simulation assumes the thermal load demand introduced by setpoint adjustments can be fulfilled by the HVAC equipment. This is a fair and necessary requirement because the setpoint schedule defined in the simulation model is typical for normal household operations, and the perturbation and optimization steps need to be informed by load changes from setpoint changes observed in the simulation. The simulation period is three months from June to August. Compared with other numerical models that use overly simplified assumptions such as a decoupled system and exclusion of solar and non-HVAC loads, this testbed is setup to represent load response authentically to setpoint adjustments without running simulation software¹.

The establishment of baseline loads, perturbations, and optimization verification are all conducted using this testbed. Figure 3-2 shows simulation results of hourly loads for 10 individual zones, with two baseline cooling setpoint schedules at 24 °C and 25 °C (with a deadband of ± 0.5 °C) for a one-week period in the summer season. The HVAC loads peak at different time of the day because of the zone orientation. HVAC loads also vary with outside

¹For the source code of the simulation, please refer to Appendix D.




Figure 3-1: A basic unit in simulation test bed, consisting of five zones (not drawn to scale)

Table 3.1: Model parameters for the simulation testbed

Parameters	Unit	Value Range
Thermal resistance (R)	$^{\circ}\text{C}/\text{kW}$	2.2 – 2.8
Thermal capacitance (C)	$\text{kWh}/^{\circ}\text{C}$	3.5 – 4.5
Initial temperature (T_{init})	$^{\circ}\text{C}$	22 – 27
HVAC capacity (P_{max})	kW	Determined by the testbed
HVAC coefficient of performance (COP)	-	2.8 – 3.3
HVAC sizing factor (sf)	-	1.4 – 2.0
Window size	m^2	{4, 5, 8}
Window solar heat gain coefficient (shgc)	-	0.3 – 0.6
Zone orientation	-	{east, west, south, north, interior}
Simulation season	-	{heating, cooling}
Zone setpoint schedule (Tset_schedule)	$^{\circ}\text{C}$	Pre-set by users or updated with the perturbation optimization outputs
Zone non-HVAC loads (Qall_schedule)	kW	Pre-set by users

temperature, with zero loads occurring at midnight as the ambient temperature drops.

Figure 3-3 illustrates room temperature dynamics with setpoint perturbation implementation for a selected zone for a one-week period with 25 $^{\circ}\text{C}$ setpoint baseline. The perturbation ± 1 $^{\circ}\text{C}$ schedule is generated from a uniform probability distribution. The room temperature follows the trend of setpoint most time except during the nighttime when room temperature drops significantly below the setpoint perturbation range due to a reduced outdoor temperature and low occupant activity level, and no solar exposure.

In this chapter, a prediction model is not applied to the simulation aggregate baseline, so the optimized load can be directly compared with the true simulated benchmark. The perturbed aggregate load is fed into a linear machine learning model. The optimization model embedding the learned perturbation model then gives optimal control setpoints. Finally,




Figure 3-2: Individual unit HVAC load for 10 simulated zones for a selected one-week period, with outdoor temperature




Figure 3-3: Room temperature dynamics with setpoint perturbation schedule for a selected zone

the optimized setpoint schedule is piped into the testbed again for validation.

3.4 Results and Discussions

3.4.1 Learning a Perturbation Model with Linear and Non-linear Algorithms

Linear Model

The perturbed dataset is split into training, validation, and test sub-sets to first select the best coefficients and tune the parameters, temperature threshold (θ_{\min}), m (memory). The best tuned parameters for this dataset are 23.3 °C (74 °F) and 24-hour memory. The selection of the parameters is based on the parameters associated with the lowest RMSE

metric in the validation dataset (Figure 3-4).




Figure 3-4: Model parameter tuning for temperature threshold (a) and memory (b) on training and validation sets

With the tuned hyperparameters, the training RMSE is 2.63 kW, and the validation RMSE is 2.57 kW. While normally the validation error is higher than the training error, the fact the training and validation datasets are not entirely stationary makes the result a little skewed. Because this linear model does not address the inherent seasonality within the dataset (i.e. August in Austin is usually hotter than May), the average ambient temperature and perturbed load are different for the two datasets.

Figure 3-5 presents the predicted and actual perturbed aggregate loads on the validation dataset. The perturbed load ranges from 10 kW to 15kW for a total of 10 zones. The predicted loads follow the actual trend, despite the hours with zero loads. The major prediction error comes from failing to capture (1) the extreme perturbed loads and (2) zero perturbed loads. For the first source of error, it is common that linear models tend to smoothen the predictions due to the averaging effect. The second source of error is also inevitable given the complexity of thermal dynamics associated with perturbation schedule. For example, a zero load could occur not only when a $+1$ °C setpoint change was implemented, but also when no setpoint adjustment occurred but the outdoor weather was cooler, or even when there were continuous setpoint perturbations in the same direction. Thus, a simple linear model would not be able to capture all these scenarios.

The inaccuracies in prediction could lead to mixed effects on the following optimization step. For example, in the cooling season, an upward setpoint adjustment decreases the total load; if the magnitude of perturbation is underestimated by the linear model, the optimization algorithm will make conservative decisions, meaning that the actual shifts in loads are greater than the predicted ones, which favors the peak reduction. On the other hand, an underestimate in perturbed loads associated with a downward setpoint adjustment may have




Figure 3-5: Predicted and actual perturbed loads for the validation dataset

the system over-respond to the setpoint change, and therefore create a new undesired peak load during the previous non-peak hours.

Non-linear Model

The non-linear multi-layer perceptron (a simple neural network model) predicts perturbed load for a 4-hour horizon, so each hourly prediction is associated with its own prediction accuracy. Figure 3-6 demonstrates the prediction RMSEs for each hourly prediction and compares them against the memory span. Intuitively, the first-hour prediction has the highest accuracy because it is closest to the historical information. On the other hand, the difference of prediction accuracy among the second to the fourth hours is insignificant. Based on the importance of each hourly prediction to the optimization problem, weights could be assigned to each RMSE to obtain the memory that fits the model best. In this case, we prioritize the prediction accuracy of the first hour to the rest, so any memory ranging from 5 to 16 hours is a good candidate. We select a 12-hour memory because this is most compatible with diurnal load patterns in reality.

Figure 3-7 shows the predicted and actual first-hour perturbed load on the validation set, with a 12-hour memory. The first-hour prediction RMSE is approximately 1.12 kW, and the rest around 1.23 kW. Compared with the linear model result, this is a 55% increase in accuracy. The improvement comes from the model's capability to capture extreme values and more precisely predict zero loads. A more accurate perturbation prediction model could prevent under- or over-estimation of thermal response and better inform the optimization algorithm.




Figure 3-6: Parameter tuning for model memory on the validation set with the neural network model




Figure 3-7: Predicted and actual perturbed load on the validation set from the neural network model

3.4.2 Setpoint Optimization and Validation

The optimized setpoint perturbation (δ_T) is used as an input to the simulation testbed to validate its actual impact on load changes, compared to the optimized load condition provided by the optimization algorithm. As is illustrated in Figure 3-8, the optimal δ_P refers to the predicted load perturbation from the optimization algorithm, whereas validated δ_P is the observed response from the testbed. It is great to see the predicted and observed load perturbations have an aligned trend. Nonetheless, there are times when loads are under- and over- predicted by the optimization algorithm compared to the validated response.

The discrepancy between predicted and observed load perturbation causes the instability of algorithm performance. Out of seven simulated days shown in Figure 3-9a, six days demonstrated varying degrees of reduction in the validated aggregate peak load demand,





Figure 3-8: Validated and optimized perturbed loads with their setpoint perturbation for a selected week

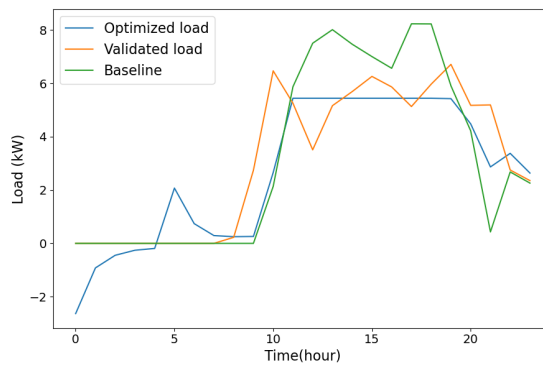
ranging from 10% to 50%. There is one day when the peak load shifts to the beginning of the day after optimization because the algorithm considers pre-cooling the spaces to minimize the peak load during the daytime. While this is not ideal, the algorithm still accomplishes load shifting to off-peak hours when grid stress is less intense. In the meantime, we also recognize that the memory of a perturbation prediction model affects the flexibility of optimization. With more memory of variables embedded in the algorithm, the optimization can control and adjust the setpoints for a wider range of hours, and therefore obtains more potential to shift and shave peak loads. Finally, besides the perturbation prediction error, the nature of a 1R1C simulation model itself could contribute to the instability of the method. Because the 1R1C model treats the entire space as one capacitor and one resistor, the thermal dynamics of air and other building materials are not separately considered. This could lead to an over-reaction of the system as both the air and building materials are modeled to be heated up and cooled down at the same rate, which is not true. In reality, building materials have much larger thermal capacitance and resistance than air, which can smooth load perturbations. To address this issue, my colleague Stella Zhang has adopted a similar three-step approach with simulations conducted in EnergyPlus that involves more realistic modeling scenarios [86]. It presents a more predictable and consistent algorithm performance than using a numerical model, albeit with more modeling efforts.

3.5 Future Work

To ensure a consistent optimization outcome, a 2R1C or 2R2C model could be developed to more realistically present thermal response of space. In the meantime, realizing a non-linear model can significantly out-perform a linear model motivates the future MIT research team to further explore a new non-linear optimization approach. Furthermore, while this



(a)



(b)

Figure 3-9: Baseline, optimized, and validated aggregate loads for a selected week (a), and a selected day (b) using the simulation testbed

three-step approach assumes a uniform setpoint perturbation across all zones, it is difficult to reach a consensus among space users. Thus, the ability to perturb setpoints on the room level would provide more flexibility to meet different space thermal comfort requirements. This may involve more challenges for baseline and perturbation models to pick up room-level signals. Finally, it is also important to learn the data requirement to train a baseline and perturbation model such that building owners and grid utilities could recognize the capital investment and implementation timeline before adopting this new technology. This problem is further elaborated in Chapter 9.

Chapter 4

Two-step GHG Emission Minimization

4.1 Background

This chapter introduces the two-step simulation with optimization approach with EnergyPlus as a room-level framework. The two-step approach is a variant of the three-step approach discussed in Chapter 3 by combining the first two steps into one perturbed load model. This simulation-based solution is then later implemented onsite experiment on MIT campus, presented in Chapters 6-8.

EnergyPlus, a whole building energy simulation program developed by the U.S. Department of Energy [28], has been considered the state-of-art simulation software since its launch in 2000. Researchers have used EnergyPlus mainly for building design [68], building load simulation [38], energy consumption analysis [32], and more recently HVAC equipment control for efficient energy performance [87, 88, 34]. The high-order simulation algorithm programmed and executed in the software allows researchers to obtain realistic simulation results compared to lower-order numerical methods. Nonetheless, establishing and running a fully parameterized model can be computationally expensive and labor intensive. To integrate a HVAC control algorithm for equipment performance improvement in a simulation model can be even more difficult as this process requires an iterative simulation process. Therefore, the two-step approach decouples the simulation from optimization to minimize the engineering effort and mitigate the required knowledge about the software deployment pipeline involved using machine learning surrogate models.

The two-step approach relies on the ability of building temperature setpoint adjustments to lower energy consumption and GHG emission. To capture and generalize thermal response of a given space, we train machine learning surrogate models on the simulated datasets. The application of machine learning methods to commercial building heating and cooling

systems is a well-studied field. These applications are broadly categorized as (a) building energy modeling, (b) optimizing energy use [69], and (c) occupancy prediction [36]. This project spans all three areas, with the first two being the focus of this chapter and the third investigated by Julia Wang, my research team member. The primary contribution of the project is experimental validation in a real building, which we have instrumented with energy sub-meters, people counters, and other sensors. A secondary contribution is optimizing for GHG emissions, rather than more common performance criteria such as energy consumption and operating costs. In [36], both experimental validation and new performance criteria are identified as important areas for future research.

4.2 Methodology

The approach involves two steps, summarized in Figure 4-1: (1) generate and learn a load profile with a perturbed temperature setpoint schedule, and (2) optimize for reduction on total GHG emission based on the perturbed load profile by constraining the setpoint temperature schedule. Specifically, we use EnergyPlus to simulate perturbed load conditions, with a random perturbation of either ± 5 °C or 0 °C and unperturbed baseline setpoint at 20 °C. This gives helpful information on load behavior under steady-state and transient state room conditions. We then apply linear machine learning algorithms to predict the load profile, using simulated hourly load data for training and validation. Finally we optimize for minimal GHG emission with a convex optimization algorithm that produces the associated hourly setpoint schedule to be fed into EnergyPlus for validation, and implemented in onsite experiments discussed in Chapters 6-8.

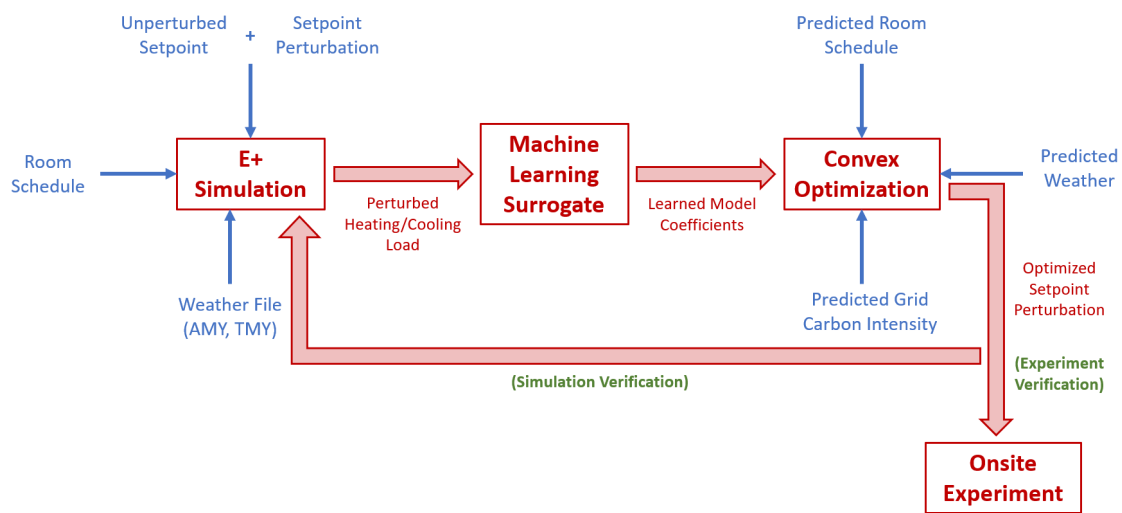


Figure 4-1: Optimized 24-hour temperature setpoint prediction system

4.2.1 EnergyPlus

EnergyPlus is employed in the design of the classroom model. The case of a medium-sized MIT classroom, in Cambridge, MA has been used for this purpose. While there are other energy simulation programs, EnergyPlus makes it possible to model new and complex building technologies which cannot be modeled by other whole building energy simulation programs. Also, EnergyPlus models can be applied to a large number of zones and whole buildings that could be used for scaling up the results for the entire campus.

The current EnergyPlus model is designed according to the real classroom parameters to the best of our knowledge. The corresponding modeled outputs provided information about the load behavior under the steady-state and transient state room conditions. The methodology used in the project has the following steps: (1) data collection about the classroom, (2) preparation of schedules for occupancy and heating/cooling using hypothetical or actual data (if available), (3) classroom model development for energy simulation model in EnergyPlus, and (4) modification of weather data file required for simulation using local Typical Meteorological Year(TMY) and Actual Meteorological(AMY) files.

Parameters used in the simulation include: general geometry of the classroom, classroom internal loads based on a pre-determined occupancy schedule and typical teaching equipment in classrooms, building envelope properties such as glazing area and materials, daylighting and solar, electrical systems and equipment, HVAC systems, occupancy (currently assumed no occupancy due to pandemic), and temperature setpoint schedule. The current setpoint profile includes a range of perturbed setpoints, for example, ± 5 °C. This classroom model simulates and outputs room loads on an hourly basis for a 3-year period¹.

The EnergyPlus model outputs are used as inputs to machine learning models. The yearly simulation is split into two modes, cooling season and heating season, to facilitate the following heating and cooling load prediction for machine learning models. For example, only the heating mode is enabled by setting a low cooling setpoint during the heating season and vice versa for the cooling season. In this way, only the heating load is generated during the heating season and cooling load for the cooling season, so that the machine learning model has a clear expectation of the prediction target during a given day.

However, this simplified method may not accurately reflect the actual room energy consumption. It is likely that space can incur both cooling and heating loads in heating or cooling seasons, depending on how room temperature is controlled. In the meantime, the arbitrarily defined heating and cooling seasons as the simulation period neglect the existence of shoulder seasons that have more interchanging heating and cooling loads throughout a day. This may cause difficulties for a learning model afterwards to pick up the load trends especially in the shoulder season.

¹Model details are included in Appendix E.

4.2.2 Machine Learning Model Exploration

This is a time-series learning task. We have experimented with several state-of-art models, such as ARIMA, Prophet model from Facebook [79], multi-layer perceptron, and recurrent neural network. To train and tune and validate model performance, we split data into training, validation and test sets. For this time series training task, each set consists of a dataset of one-year continuous simulation results. Two models are individually trained for heating and cooling loads.

Model Features and Target

The forecasting target is the hourly heating or cooling load.

It is common knowledge that the more modeling features provided to a learning task, the better the outcome would be. And indeed, the simulation output includes a large number of modeling parameters that are important to load forecasting. However, we only select the following as features considering their availability in practice: historical load condition (P), ambient dry-bulb temperature (T_{out}), room setpoint (T_{set}), change in temperature setpoint with respect to the previous time step ($dT_{set}(t) = T_{set}(t) - T_{set}(t - 1)$), temporal embedding. Specifically, we applied a sine and cosine embedding to all temporal information, as is discussed in [40] to capture daily, weekly, and seasonal trends of rooms load condition influenced by occupant behaviors, building schedules, and seasonality. For example, to represent "time of day", we used $[\sin(\pi(h)/12), \cos(\pi(h)/12)]$ to represent a 24 hour cyclic nature explicitly in the learning problem. In the case of the simulation dataset, we encode only time of day as a temporal feature. Additionally, to ensure model convergence, we scale T_{out} , T_{set} , and P to standard normal distributions.

Teacher Forcing Training for Time Series Inputs

Teacher forcing is a widely used technique developed several decades ago for training sequential input data. It was originally developed for quickly and efficiently training recurrent neural network models, used in the field of natural language processing, and later was extended to applications in other fields. The time series prediction task requires the prior time step predicted room load to become an input to the next step load prediction. When teacher forcing is enabled, the model uses ground truth as input (as if there was a teacher telling the right answer), instead of the output predicted in the prior step, in the training phase. The two training methods (Autoregressive with exogenous regressors and Multi-layer perceptron) presented in this project also adopt this technique.

Autoregressive with Exogenous Regressors (ARX)

A wide range of statistical methods have been developed for time series analysis, for example, Autoregression (AR) that predicts future behaviors only with past behaviors as inputs, Autoregressive Moving Average (ARMA) that considers the observations and residual errors, Autoregressive Integrated Moving Average (ARIMA) that takes a differencing pre-processing step of the sequence to make it stationary on top of ARMA model, and Seasonal ARIMAX that considers seasonal trends and exogenous variables.

All the model variations have been explored. While domain knowledge may assert a strong seasonality and trend in the thermal load, we find that the extra model complexity introduced by "Integration," "Moving Average" or "Seasonality" terms only marginally improves model performance. Considering the trade-off between model complexity and performance, we selected ARX(6, 1) as the preferred model. The autoregressive component includes the previous 6-hour load which has a lag of 6, and the exogenous regressors consist of current T_{out} , T_{set} , dT_{set} , and time of day which has a lag of 1.

Mathematically, an ARX(6, 1) model is defined as:

$$y(k+1) = a_1y(k) + a_2y(k-1) + \dots + a_6y(k-5) + \mathbf{b}_1\mathbf{U}(k) \quad (4.1)$$

where y is the thermal load, the target output, a_1, \dots, a_6 are the coefficients of the autoregressive components, \mathbf{U} is the exogenous vector variable, and \mathbf{b}_1 is the coefficient vector associated with the exogenous vector variables.

Prophet Model

ARX models are interpretable and have strong forecasting performance when the lags parameters are trained well, yet they lack in capturing the non-linearities in the time-series trends. To address those challenges, we leverage the Prophet model [79] to forecast the hourly heating load with the outdoor temperature and setpoints as regressors. The Prophet model is an adaptation of a decomposable time series model [44] with three time-series components: trend, seasonality, and holidays.

$$y(t) = g(t) + s(t) + h(t) + e(t) \quad (4.2)$$

where $g(t)$ is the trend function which models non-periodic changes in the value of the time-series, $s(t)$ represents periodic changes (e.g., weekly and yearly seasonality), and $h(t)$ represents the effects of holidays which occur on potentially irregular schedules over one or more days. The error term $e(t)$ represents any idiosyncratic changes which are not accommodated by the model; later we make the parametric assumption that is normally distributed. This specification is similar to a generalized additive model (GAM) [45], a

class of regression models with potentially non-linear smoothers applied to the regressors. Here we use only time as a regressor but possibly several linear and non-linear functions of time as components. Modeling seasonality as an additive component is the same approach taken by exponential smoothing [39]. Multiplicative seasonality, where the seasonal effect is a factor that multiplies $g(t)$, can be accomplished through a log transform. The GAM formulation has the advantage that it decomposes easily and accommodates new components as necessary, for instance when a new source of seasonality is identified.

We trained a Prophet model with $y(t)$ hourly heating load, and incorporated the outdoor temperature and heating setpoints as additional regressor. We find the model to be best trained when the changepoint prior scale is tuned to 0.4. Prophet includes functionality for time series cross validation to measure forecast RMSE using historical data. We specify the initial training period as 60 days, and both the forecast horizon and the spacing between cutoff dates as one day.

Multi-layer Perceptron (MLP)

Previous research efforts suggest that more complicated neural network models for sequential prediction tasks, such as Recurrent Neural Nets (RNN) and its variant Long Short-Term Memory (LSTM), are too powerful for load prediction because the training data have a very low-dimension data structure compared to language model training inputs for which those models were designed. Such complex model structures require a lot of computation resources while not generating satisfying outcomes.

Therefore, we decide to use the simplest version of a neural network, a multi-layer perception (MLP) to embed the non-linearity relationships among the variables. This MLP consists of three layers of nodes: an input layer, a hidden layer, and an output layer, accompanied by a ReLU activation function. To encode the time series properties, the model inputs and structures (training and validation) are structured in the same way as an ARX model. The MLP takes in historical 24-hour load and current T_{out} , T_{set} , dT_{set} , and time of day as inputs and forecasts the load conditions for the next 24 hours.

4.2.3 Convex Optimization

After training an energy prediction model, we embed it in an optimization algorithm to generate a temperature setpoint profile that minimizes GHG emissions. While we have tried various types of machine learning models described in the previous section, we apply a convex optimization model combined with the ARX(6, 1) prediction model. The problem of minimizing cumulative GHG emissions takes the following general form:

$$\min_{x \in \mathcal{X}} \{ \mu^\top y \mid y = f(x) \} \quad (4.3)$$




Figure 4-2: An illustration of a 3-layer MLP model

Here the decision variable $x \in \mathbb{R}^n$ ($^{\circ}\text{C}$) contains the temperature setpoints at each step of the time horizon, $\mathcal{X} \subset \mathbb{R}^n$ is a convex set of feasible temperature setpoint profiles, $\mu \in \mathbb{R}^n$ (kg/kWh) contains the GHG intensity of electricity at each time step, $y \in \mathbb{R}^n$ contains the energy used at each time step, and $f : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^m$ is the energy prediction model. (In this abstract representation, the structure of f includes the exogenous features and the model parameters.)

Note that the optimization is conducted on a room-level that ignores the complex energy supply (a tri-generation) system on campus. Additionally, while the algorithm takes grid carbon intensity as a basis for optimization, the campus plant largely consumes natural gas. Those implications are elaborated in detail in Chapter 9.

Since the selected prediction model is linear in the temperature setpoints, then the objective function in Problem (4.3) is linear in x . In this case, Problem (4.3) is convex, as the set \mathcal{X} is convex by assumption. Therefore, the problem can be solved to global optimality in polynomial time using, *e.g.*, interior-point methods. In our applications, the set \mathcal{X} can usually be described by a system of linear inequality constraints, so Problem (4.3) reduces to a linear program. It can therefore be solved efficiently and reliably by off-the-shelf software, in this case the CVX (in MATLAB)/ CVXPY (in Python) package developed by Stephen Boyd’s group at Stanford [29, 30, 14].

The optimization model outputs the next 24-hour optimal setpoint schedule based on its best knowledge of the room and external conditions for the next day. It relies on the predicted or known information provided by the forecasting models. For example, the grid intensity forecast tells how clean tomorrow’s grid would become, and the room occupancy model shows if the space is occupied during a certain period. In this study, we select one-

day grid intensity historical 5-minute data of ISO New England from ElectricityMap² as predetermined values to make the results more interpretable .

The following two sets of constraints are implemented in the optimization process. First, setpoint range. Take the heating season model as an example, heating setpoints range from 15 to 24 °C; when room is predicted as unoccupied, the heating setpoint should be kept between 20 and 21 °C; otherwise, the range can be extended to 15 to 24 °C to allow for more natural temperature drifts. Second, the rate of change of setpoints. Considering the heating and cooling capacity from real HVAC equipment, we impose the maximal hourly setpoint change rate as 0.5 °C. For heating season, this means the next time step setpoint cannot be more than 0.5 °C higher than the current one (no constraint on the downward change rate). This ensures the optimal setpoints implemented in experiments can be later closely tracked by the HVAC equipment.

This optimization could be further modified to incorporate occupants' thermal comfort as the second objective in addition to GHG emission. Then the problem becomes a multi-objective optimization framework, and depending on the importance of the two objectives, weights are assigned to each objective respectively. To quantify thermal comfort, we could introduce a quadratic penalty term that measures the deviation of room temperature from the desired temperature (preferred by occupants). The multi-objective problem will be discussed in detail in Chapter 5. The following optimization results are all based on the setup of a single objective problem³.

4.2.4 EnergyPlus Simulation Verification

To evaluate the GHG savings from the optimal setpoint schedule, we pipe the 24-hour schedule back to EnergyPlus using eppy [66], a Python based package to modify the simulation. Then the optimized GHG is compared with the baseline GHG.

4.3 Results and Discussion

4.3.1 Machine Learning Model Performance Comparison

To evaluate and compare the model performance, we feed the model with the test set of heating loads and compare the average Root Mean Square Error (RMSE), listed in Table 4.1. A 24-hour forecasting horizon of the three models from a selected period is shown in the Figure 4-3. Overall, while not perfect, all three models correctly forecast general load trend. Compared to Prophet and ARX, MLP can better capture extreme load conditions such as peaks. ARX and Prophet share very close RMSE, around 4MJ, while MLP significantly outperforms both, giving only 1.8MJ RMSE.

²<https://app.electricitymap.org/map>

³All code documentation is posted on the GitHub Repo; link is provided in Appendix D.

Table 4.1: Model performance comparison

	ARX(6,1)	Prophet	MLP
RMSE [MJ]	4.25	4.04	1.77




Figure 4-3: Model performance comparison on a 24-hour forecasting horizon from the test set. Black solid line represents true load, and the dashed orange, red and blue lines are from Prophet, ARX, and MLP models respectively.


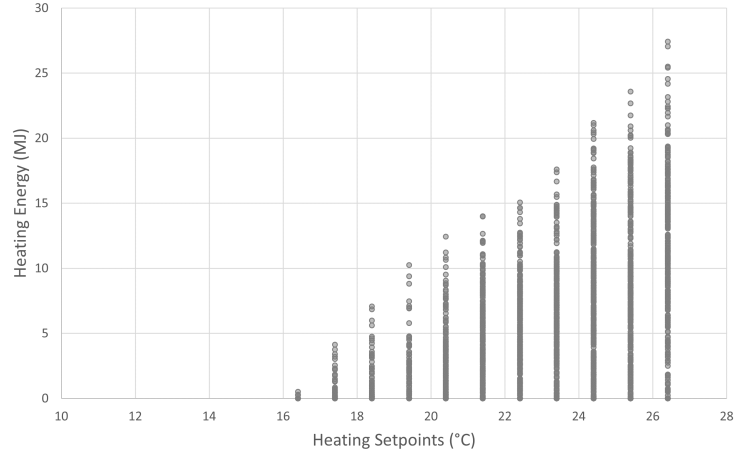



Figure 4-4: Actual and predicted (with ARX model) hourly heating load trend in the heating season

More complex non-linear neural networks are often expected to outperform simpler linear models. Carefully investigating the simulation dataset, we recognize that the trends within the data are inherently nonlinear. Figure 4-4 illustrates the predicted and actual hourly heating load in the heating season (Oct. to Mar.) While it is expected the heating loads would peak in Jan. that is usually associated with lowest outdoor temperature, we only observe some seasonality in the dataset. This is determined by how the dataset is generated.



(a)



(b)

Figure 4-5: Hourly heating energy associated with heating setpoints (a) and outdoor temperature (b)

Recall that the first phase in this two-step approach, room setpoints are perturbed by ± 5 $^{\circ}\text{C}$. This means the setpoint difference between two hours can range from 0 $^{\circ}\text{C}$ (no change applied to the previous hour) to 10 $^{\circ}\text{C}$ (-5 followed by $+5$ or vice versa). The setpoint perturbations disrupt the linear or quadratic relationship between setpoints and heating load, as described in Figure 4-5a for a simulation in the heating season. While a higher heating setpoint is associated with larger heating energy, it is still possible that a higher setpoint can lead to zero load and a lower setpoint to a non-zero load because the load then is driven by setpoint changes rather than setpoints themselves. In addition, the direct relationship between ambient temperature and heating load becomes unclear, shown in Figure 4-5b. The room could have zero load when ambient temperature is very low but the heating setpoint just gets adjusted downward by 5 $^{\circ}\text{C}$ (such as the zero-load data at -10 $^{\circ}\text{C}$ outdoor temperature in Figure 4-5b). All the non-linear properties introduced by setpoint

perturbations make a non-linear model a more suitable candidate for the prediction task.


4.3.2 Optimization Results

During the optimization process, a hypothetical room occupancy schedule is fed into the framework, along with a grid marginal operating emission rate (MOER) as an indicator for grid carbon emission intensity. In this simulation, the room occupancy is labeled either as 0 (unoccupied) or 1 (occupied). The MOER schedule comes from a real 24-hour period from ISO New England obtained from ElectricityMap. The optimization randomly selected a heating day using the 2017 Boston AMY weather file as the ambient temperature.


After performing sensitivity analysis, we find that the algorithm is most sensitive to the occupancy schedule, through the following three test trials. In each trial, we control all other variables except for the occupancy schedule and observed the optimal setpoint schedule. Each trial result is presented in two figures, with the top one showing the setpoints and grid MOER and bottom one for setpoints and occupancy status.

The first test shown in Figure 4-6 examines the algorithm’s response to occupancy changes in the form of a delta function. Aware of the occupancy schedule change, the algorithm mandates the setpoint during the occupied period above 20 °C. To achieve this, considering the constraint on ramping rate, the setpoint gradually increases starting from midnight. This setpoint profile also validates the correctness of the embedded load prediction model – it recognizes a higher heating setpoint is associated with larger room loads, and therefore adjusts the setpoints downward during the unoccupied time to minimize total energy and GHG emission. However, since we do not impose a setpoint ramp down rate, it is expected that the algorithm should mandate a drastic setpoint drop after the 10th hour when room is no longer occupied. It could have achieved a larger saving if it released setpoint to the lowest allowed 16 °C immediately.

The second test scenario describes a 24-hour full occupancy room status. It is designed to test the model’s response to varying MOER levels in the grid. A full occupancy schedule requires the all-day setpoints constantly above 20 °C. Nonetheless, because the MOER level in the specific simulation period ranges from 250 g CO₂eq/kWh to 330 gCO₂eq/kWh, we would expect a load shifting strategy that preheats the room before the grid becomes carbon intensive. Unfortunately, this is opposite to our observation. As is shown in the Figure 4-7a, the highest setpoint occurs when the grid has the highest MOER. While the magnitude of the change is very subtle, it shows that the model is less sensitive to MOER and therefore may not bring pre-conditioning strategies to achieve load shifting. This can be explained by two factors. First, the variation in ISO New England MOER is limited. The selected MOER profile has a daily variation of approximately 30% (note that 30% variation is not common in ISO New England market and typical variation may be below 20%), which is relatively small compared to that in California or Texas where solar and




(a)




(b)

Figure 4-6: (a) Optimized setpoints vs. grid MOER for test case 1; (b) Optimized setpoints vs. room occupancy status for test case 1



(a)




(b)


Figure 4-7: (a) Optimized setpoints vs. grid MOER for test case 2; (b) Optimized setpoints vs. room occupancy status for test case 2

wind renewable energy creates a more fluctuated grid condition. We envision that the algorithm would pick up the MOER signals more effectively if the simulations were set in a more fluctuated electricity market. Second, the embedded load prediction model cannot accurately forecast load conditions due to the inherent non-linearity in the load behavior, as is described in the previous result section. The optimal control depends on the accuracy of the embedded model that informs the optimization algorithm of the predicted outcome of the planned decision horizon. Consequently, a less accurate prediction model may not capture fast changing dynamics presented in the simulation, which renders it difficult to guide the following decision-making process.

The last test simulates a typical classroom occupancy condition – class starts in the early morning, followed by a noon break, and ends by late afternoon. Figures 4-8a and 4-8b demonstrate the model’s capability of handling more complicated real-life scenarios. This



(a)




(b)

Figure 4-8: (a) Optimized setpoints vs. grid MOER for test case 3; (b) Optimized setpoints vs. room occupancy status for test case 3


assures that the two-step approach generates a reasonable, even though not perfect, setpoint control scheme. Next, we want to further quantify the savings from the optimal control by piping the setpoints back to the EnergyPlus simulation and comparing the calculated GHG emissions with the baseline condition.

4.3.3 Load Minimization Verification


In the baseline model, we set a constant setpoint at 20 °C, corresponding to the setpoints designed for the occupied period. The validation model uses the optimal heating setpoints. Using the three test cases in the previous sections, we obtained the following GHG emission profiles shown in Figure 4-9, and a savings summary in Table 4.2.



(a)



(b)



(c)

Figure 4-9: Validated vs. baseline GHG emissions. (a): test case 1; (b): test case 2; (c): test case 3⁴

All three tests share the same baseline, with a total GHG emissions of 5.38 kg of CO₂ eq. throughout the 24-hour operation. All test cases manage to save emissions, from the least 2%, to the most 45%. The amount of savings in this case depends on the occupancy schedule. To illustrate, the second test case uses a fully occupied schedule which also requires

⁴Note that the room does not generate heating load during the daytime due to room plugin load and solar load, so both validated and baseline results have zero emission. The same applies to the other test cases.

Table 4.2: Baseline and validated optimal GHG emission comparison

Total GHG Emission (kg)	Test 1	Test 2	Test 3
Baseline	5.38	5.38	5.38
Validated Optimal	3.35	5.27	2.96
Percentage savings	37.7%	2.0%	45.0%

a minimum 20 °C heating setpoint throughout the day, similar to that in the baseline. Because the model is unable to provide load shifting strategies, the baseline and optimal solutions provide almost identical GHG emissions. On the other hand, the first and third test cases benefit from the schedule flexibility, which allows for a lower setpoint during the unoccupied period to save unnecessary heating loads. This result suggests that getting knowledge of the room’s occupancy status is critical to maximizing energy saving potential. However, we also need to acknowledge that while the simulation results are promising with a complete and accurate forecast of future external conditions, a perfect prediction of the future is never possible in reality. More simulation needs to be carried out to analyze the robustness of the model, that is, model performance when the forecast is different from the reality.

4.4 Future Work

While the convex optimization integrated with the linear energy prediction model gives promising results, we realize this method has its own limitations. For energy prediction, linear models are significantly less accurate than nonlinear alternatives. Such prediction inaccuracies lead to the following optimization’s insensitivity to grid carbon level, which disable the load shifting strategies that could otherwise achieve more GHG emission savings. For this reason, for future work, we are considering a two-stage approach that uses both linear and nonlinear energy prediction models. The first stage of this approach involves solving a convex version of Problem (4.3) with a linear energy prediction model $f^{(\ell)}$ to generate a solution $x^{(\ell)}$. The second stage involves solving a nonconvex version of Problem (4.3), with the linear model $f^{(\ell)}$ replaced by a nonlinear model $f^{(n)}$ in the form of neural network models, regression trees or other advanced prediction algorithms.

While there are no guarantees that the second-stage problem can be solved to global optimality, it can be solved locally by various gradient descent algorithms. These algorithms can be warm-started with the first-stage solution $x^{(\ell)}$. Given that warm-start, running gradient descent on the second-stage problem should generate a solution at least as good as $x^{(\ell)}$, and possibly significantly better. However, solving the second-stage problem will likely require coding our own local optimization routine and, in particular, computing derivatives of the nonlinear energy prediction model $f^{(n)}$. We are hoping that the proposed nonlinear

models combined with a non-convex optimization approach will be able to bring about more advanced, robust, and accurate setpoint control algorithms to be applied in a variety of settings.

Chapter 5

Model Predictive Control for A Multi-objective Optimization

5.1 Background

While the previous two chapters have explored self-defined predict-and-optimize frameworks, to minimize GHG emissions through HVAC setpoint control, this chapter seeks to adopt a state-of-art control algorithm, model predictive control (MPC), to formulate a multi-objective optimization problem that considers emission and cost factors with constraints on thermal comfort and indoor air quality.

Over the past few years, researchers have analyzed the opportunities offered by implementing classical control principles on HVAC. These studies reveal that the MPC technique has the potential of improving building energy efficiency. Ma et al. [54] showed that MPC achieves 19% energy savings in a university campus building in Berkeley. Ma et al. [53] analyzed the economic saving potentials with MPC in the range of 16-26%. The key difference between MPC and traditional control algorithms is MPC's predictive capacity. Its capability of making decisions based on predicted future scenarios gives considerable power of planning an optimized control schedule for the near future. MPC shows strength in building HVAC operation control because it can optimize building energy consumption, cost, and GHG emission based on the predicted conditions in the electricity market. For example, MPC can leverage opportunities such as changes in grid carbon emission intensity to strategically manipulate the HVAC operation schedule, which leads to minimum operation cost and GHG emission.

Although recent research has extensively investigated MPC used for building HVAC system control to reduce energy consumption, not many groups have explored its application in a multi-objective building performance optimization problem. Previous research focuses on minimizing HVAC operation carbon footprint or energy cost, but ignores the importance

of indoor air quality (IAQ), which plays a crucial role in indoor environmental quality (IEQ). A common industry practice of maintaining IAQ is by providing the minimum amount of outdoor air required by ASHRAE 62.1 Standard [18]. However, such a minimum ventilation rate may not be sufficient to maintain a healthy IAQ especially when room occupancy is higher than the designed occupancy. Hence, the goal of this project is to establish a multi-objective optimization framework that minimizes GHG emission and cost associated building HVAC system operation, while maintaining thermal comfort (defined by indoor temperature range) and IAQ (defined by room CO₂ level), based on the knowledge of occupancy schedule and outdoor weather forecast.

5.2 Methods

To model dynamical behaviors of indoor CO₂ concentration and room temperature, system power consumption, total electricity cost, and GHG emission, we need a combined optimization and simulation model that includes the following components.

5.2.1 Thermal Dynamics: Resistance Capacitance (RC) Network Model

The optimization is embedded with a 1R1C network model, with one capacitor, equivalent to thermal capacitance of mass and air, and one resistor that lumps all heat transfer effect such as conduction and convection into one process defined in Equation (5.1). Here, T_a represents ambient temperature; T is indoor air temperature that is assumed the same as mass temperature; thermal power \dot{Q}_{hvac} is heat injected by controlled thermal equipment, such as HVAC equipment; and thermal power \dot{Q}_e is heat injected by exogenous sources that contribute to thermal disturbance, which may include solar radiation, heat brought by outdoor air, and occupant and lighting load. In this model, we do not account for solar radiation. Thermal capacitance and resistance (R and C) are properties of room building material.

$$\dot{Q}_c + \dot{Q}_e = C \frac{dT}{dt} + \frac{T - T_a}{R} \quad (5.1)$$

The solution of this first-order model gives a linear representation of the system

$$T_r(k+1) = aT_r(k) + (1-a)[T_{out}(k) + R(\dot{Q}_{hvac}(k) + \dot{Q}_e(k))] \quad (5.2)$$

where

$$a = e^{-\frac{\Delta t}{RC}}$$

$$\dot{Q}_{hvac}(k) = \dot{m}(k)C_{p,air}(T_r(k) - T_{sup}(k))$$

\dot{Q}_{hvac} is calculated by multiplying recirculated air mass flowrate \dot{m} by specific heat capacity of air $C_{p,air}$, and temperature difference between room T_r and supply air T_{sup} ,

Table 5.1: Model variables, parameters, constraints, and objectives

Variables	Symbol	Unit	Description	Upper/Lower Bounds	Nominal Value	
Dependent Variables	HVAC energy flow	Q_{hvac}	Heat injection/extraction rate	Equipment capacity	-	
	Volumetric flow rate	\dot{V}	Outdoor air exchange rate	Equipment capacity	-	
	Room temperature	T_r	Room temperature	Defined by constraints	-	
	HVAC air flowrate	\dot{m}	Mass flowrate of HVAC recirculated air	Equipment capacity	-	
	Room CO ₂ level	C_r	CO ₂ concentration	Defined by constraints	-	
	Air change rate	λ	Defined as $\frac{\dot{V}}{V}$	-	-	
	Outdoor air heat flow	Q_{out}	Heat flow brought in by outdoor air	-	-	
	HVAC power	P_{hvac}	HVAC power consumption	Equipment capacity	-	
	Fan power	P_{fan}	Fan power consumption	-	-	
	Plant power	P_{plant}	Heating/cooling plant power	-	-	
	DOAS power	P_{doas}	DOAS power consumption	-	-	
	Normalized power	H	Normalized DOAS fan power	Equipment capacity	-	
	Normalized air flow	f	Normalized air volumetric flowrate	0-1	-	
	Parameters	Room thermal resistance	R	Room material properties	-	0.01
		Room thermal capacitance	C	Room material properties	-	1.8×10^6
Outdoor temperature		T_{out}, T_a	Temperature from weather forecast	-	-	
Exogenous heat flow		\dot{Q}_e	Heat injected by occupants & lighting	-	-	
# of occupants		N	Number of occupants	Defined by schedule	0-15	
Supply air temperature		T_{sup}	Constant HVAC supply air temperature	-	15	
Air heat capacity		$C_{p,air}$	Specific heat capacity	-	1000	
Room volume		V	room volume	-	120	
Room area		A	room area	-	40	
CO ₂ emission rate		\dot{E}	CO ₂ emission rate per person	-	42	
Outdoor CO ₂ level		C_{out}	Constant outdoor CO ₂ level	-	0.74 (410 ppm)	
Target temperature		T_t	Used to calculate P_t	-	-	
Target CO ₂ level		C_t	Used to calculate P_{CO_2}	-	-	
Air density		ρ_{air}	Air density	-	-	
Design fan power		P_r	Reference DOAS fan power	-	1.2	
Design air flowrate		\dot{V}_r	Reference DOAS air flowrate	-	1000	
Carbon intensity		MOER	marginal operating emission rate	-	0.94	
Electricity price		LMP	Locational marginal price	Varies with time	-	
Constraints		$T_r >= T_{r,low}$	T_r, low	Lower bound room temp setpoint	Depends on #occupants	15-20
		$T_r <= T_{r,high}$	$T_r, high$	Upper bound room temp setpoint	Depends on #occupants	27-32
		$\dot{V} >= \dot{V}_{min}$	\dot{V}_{min}	Outdoor air minimum flowrate	Depends on #occupants	16.6/ppl, 2.2/ m^2
		$C_r <= C_{r,h}$	$C_{r,h}$	Room CO ₂ Limit	Depend on #occupants	2.7-3.6
		Cost	Pr	Energy cost	$>=0$	-
Objective	CO ₂ emission	E_c	Emission from building operation	$>=0$	-	
	Thermal penalty	P_t	Thermal comfort penalty	-	-	
	IAQ penalty	P_{CO_2}	Indoor air quality penalty	-	-	

^aEach model variable will be explained again in the model context.

^bSome upper and lower bounds are not explicitly defined because they vary based on room occupancy condition.

^cParameters such as outdoor temperature do not have nominal values because they vary with time.

which is held constant at 15 °C. Δt is discrete time interval. \dot{Q}_{hvac} is positive in the heating season and negative in the cooling season.

5.2.2 Indoor Air Quality(IAQ): CO₂ Mass Balance Model

CO₂ concentration is a commonly-used indicator of IAQ. Outdoor air CO₂ level is set at 410 ppm, whereas indoor level increases with the presence of occupants who exhale CO₂. A higher ventilation rate, or a higher exchange rate of outdoor air with indoor air, lowers indoor CO₂ level. Therefore, room CO₂ dynamics can be modeled as a mass balance [23]:

$$V \frac{dC_r}{dt} = \dot{V}C_{out} - \dot{V}C_r + \dot{E} \quad (5.3)$$

where V is volume of room, \dot{V} is volumetric flowrate of outdoor air into room, C_r is room CO₂ concentration, and \dot{E} is CO₂ emission rate by occupants.

The solution to Equation (5.3) is

$$C_r(k+1) = bC_r(k) + (1-b) \left[C_{out} + \frac{\dot{E}(k)}{\dot{V}(k)} \right] \quad (5.4)$$

where

$$b = e^{-\lambda \Delta t}$$

and λ is Air Change Rate (1/hr), calculated as outdoor air volumetric flowrate divided by room volume (Equation (5.5)), which represents the frequency at which indoor air gets completely replaced by outdoor air in an hour.

$$\lambda = \frac{\dot{V}}{V} \quad (5.5)$$

We treat outdoor air intake as a separate module from the heating and cooling module to avoid non-convexity. A Dedicated Outdoor Air System (DOAS) is used to fulfill this task. As this DOAS system is not equipped with heat recovery capacity, it brings heat flow \dot{Q}_{out} (part of \dot{Q}_e) into room along with outdoor air:

$$\dot{Q}_{out}(k) = \rho_{air} \dot{V}(k) C_{p,air} (T_{out}(k) - T_r(k)) \quad (5.6)$$

5.2.3 HVAC Power Model

An HVAC power model can be decomposed into a fan power model and a heating/cooling plant power model.

$$P_{hvac}(k) = P_{fan}(k) + P_{plant}(k) \quad (5.7)$$

A heating/cooling plant model further depends on multiple variables such as heat exchange in coils, reference heat exchange rate, operation temperature, and chill water supply; a fan power model depends on parameters such as reference air mass flowrate, static pressure and maximum designed fan power. We apply a surrogate model developed in [24] which generalizes the major HVAC power consumption behavior as a function of outdoor temperature T_{out} and equipment load \dot{Q}_{hvac} :

$$P_{hvac}(k) = \left(C + \frac{1}{cT_{out}(k) + d} \right) Q_{hvac}(k) \quad (5.8)$$

where¹ $C = 0.1214$, $c = -0.1741$, and $d = 8.3356$.

5.2.4 Dedicated Outdoor Air System (DOAS) Power Model

We model a DOAS equipped with a variable speed drive without heat recovery. As proposed by Englander and Norford [33], the normalized fan power H is a function of normalized flow f described in the following equation [24]:

$$H(k) = 1.2842f^3 - 1.3156f^2 + 10.47f \quad (5.9)$$

where

$$H(k) = \frac{P_{doas}}{P_r}$$

$$f(k) = \frac{\dot{V}(k)}{\dot{V}_r}$$

We set 1000 W as the DOAS reference fan power P_r and $0.94 \text{ m}^3/\text{s}$ as reference flowrate \dot{V}_r .

5.2.5 Grid Carbon Intensity

When optimizing total CO₂ emission rate, information about local electricity grid carbon intensity is needed. Marginal operating emission rate (MOER) serves as a widely used indicator of grid carbon intensity. Companies such as WattTime use predictive machine learning algorithms to extrapolate future grid MOER based on grid power demand and supply forecast. MOER becomes higher when power demand surpasses power supply. As regulations and policies are developed to address climate change, the electricity grid carbon intensity has declined in recent years. However, the introduction of renewable energy sources has made the supply side less stable, which incurs seasonal and daily MOER fluctuations. In this model, we use historical MOER data in Massachusetts provided by WattTime [10], sampled at a 5-min interval to calculate CO₂ emission.

¹We model P_{fan} in HVAC only as a function of Q_{hvac} , and P_{plant} as a function of T_{out} and Q_{hvac} .

$$E_c(k) = MOER(k)(P_{hvac}(k) + P_{doas}(k)) \quad (5.10)$$

5.2.6 Electricity Wholesales Price

The pricing mechanism works similarly as the grid emission intensity pattern. When demand surpasses supply, electricity generated at a higher cost is brought into the market. Locational marginal pricing (LMP) [4] is representative of electricity wholesale price. Market operators utilize time-series algorithms to predict day-ahead LMPs with knowledge about day-ahead power demand and supply. In this model, we use real-time LMP data from ISO New England, sampled at a one-hour interval, to calculate energy cost.

$$Pr(k) = LMP(k)(P_{hvac}(k) + P_{doas}(k)) \quad (5.11)$$

5.2.7 Occupancy Schedule

We hypothetically define a room schedule that resembles a typical campus room schedule. In the case of MIT classrooms, room schedule and occupancy rate are based on course registration information. Given occupancy status, we can define room a setpoint temperature range, room CO₂ concentration limit, CO₂ emission rate from occupants, heat injected by occupants, and minimum required ventilation rate following ASHRAE Standard 62.1 [18].

(1) When room is scheduled as occupied at timestep k :

- Temperature setpoint lower bound: $T_{r,l}(k) = 20 \text{ }^\circ\text{C}$
- Temperature setpoint higher bound: $T_{r,h}(k) = 27 \text{ }^\circ\text{C}$
- CO₂ concentration limit: $C_{r,h} = 2.7 \text{ g/m}^3$ or 1500 ppm
- Target room temperature $T_t = 24 \text{ }^\circ\text{C}$ with a deadband of $\pm 0.5 \text{ }^\circ\text{C}$
- Target room CO₂ level $C_t = 1.8 \text{ g/m}^3$ (1000 ppm) a deadband of $\pm 0.5 \text{ g/m}^3$ (280 ppm)

A penalty will occur if the room condition goes beyond the deadbands with respect to T_t and C_t .

(2) When room is scheduled as unoccupied at timestep k :

- Temperature setpoint lower bound: $T_{r,l}(k) = 15 \text{ }^\circ\text{C}$
- Temperature setpoint higher bound: $T_{r,h}(k) = 32 \text{ }^\circ\text{C}$
- CO₂ concentration limit: $C_{r,h} = 3.6 \text{ g/m}^3$ or 2000 ppm

Target room temperature and CO₂ concentration are not set for unoccupied room status; neither is a penalty function.

(3) Other parameters defined:

- Lighting load density: 10 W/m^2
- Occupant load: 100 W/pppl
- Minimum ventilation rate of air:
 $\dot{n}_{min} = 20 \text{ kg/hr/pppl}$
 $\dot{n}_{min} = 2.6 \text{ kg/hr/m}^2$
- CO₂ emission rate: $42 \text{ gCO}_2/\text{hr/pppl}$

5.2.8 Weather Forecast

The outdoor weather condition affects room thermal performance as heat transfer processes take place in the form of convection and conduction between outdoor air and room. In this model, we use historical Cambridge weather data as our hypothetical weather forecast. Only dry-bulb temperature is considered; in reality, we can easily get this information (one-day ahead forecast) through weather forecast stations.

5.2.9 Objective and Constraints

Given the K time steps for optimization, we have the following objective function:

$$\sum_1^K w_1 Pr(k) + w_2 E_c(k) + w_3 P_t(k) + w_4 P_{CO_2}(k) \quad (5.12)$$

for $k = 1, 2, 3, \dots, K$.

where w_1, w_2, w_3, w_4 are weights of objectives corresponding to price, GHG emission, thermal comfort penalty, and IAQ penalty, and

$$P_t(k) = \frac{1}{T_t^2} [\max((T_t - 0.5) - T_r(k+1), 0) + \max(T_r(k+1) - (T_t + 0.5), 0)]^2 \quad (5.13)$$

$$P_{CO_2}(k) = \frac{1}{C_t^2} [\max((C_t - 0.5) - C_r(k+1), 0) + \max(C_r(k+1) - (C_t + 0.5), 0)]^2 \quad (5.14)$$

Subject to:

$$\dot{V}(k) \geq \dot{V}_{min}(k) \quad (5.15)$$

$$T_{r,l}(k+1) \leq T_r(k+1) \leq T_{r,h}(k+1) \quad (5.16)$$

$$C_r(k+1) \leq C_{r,h}(k+1) \quad (5.17)$$

$$|\dot{Q}_{hvac}(k+1) - \dot{Q}_{hvac}(k)| \leq 4500 \quad (5.18)$$

Note that in Equations (5.13) and (5.14), both penalty values are normalized by their own magnitude. Such a penalty mechanism would be activated only when the room is occupied. Ideally, the CO₂ concentration should not have a lower bound because IAQ becomes better (assuming no outdoor air pollution) when more fresh air gets exchanged. The main purpose of setting a lower bound is mainly to prevent model instability. Equation (5.18) requires that change in energy flow between each time step should not be greater than 4500 W to prevent equipment wear and tear.

5.3 Convex Optimization with Non-convex Simulation

The optimal control decision is reached by a combined process of simulation and optimization. As defined in the previous section, each module simulates future room condition $C_r(k+1)$ and $T_r(k+1)$ and estimates future power consumption $P(k+1)$ and carbon emission $E_c(k+1)$ based on current room condition $C_r(k)$ and $T_r(k)$, along with scheduled occupant activities and outdoor weather forecast. The algorithm tries to find an optimal and executable control solution at the current stage that fulfills all constraints on future room conditions. Therefore, the MPC algorithm requires certain knowledge about future events. If limited information is available, the model will act in response to the near future, similar to a greedy optimization algorithm; on the other hand, if sufficient and accurate information is provided, the model will make more informed decisions and show a greater level of foresight.

Because the model uses discretized solutions to dynamical systems for modeling room thermal behavior and mass transfer, a small simulation time step is required to capture fast-changing room dynamics. Specifically, we simulate and optimize on a five-minute interval in an alternating fashion.

Algorithm 1: MODEL PREDICTIVE CONTROL (MPC)

input: Initial room condition $[c_r, t_r]$, total step n , outlook K
init $[\mathbf{C}_r, \mathbf{T}_r]$
 $[\mathbf{C}_r(1), \mathbf{T}_r(1)] \leftarrow c_r, t_r$
for $i \leftarrow 1$ **to** n **do**
 $[\dot{\mathbf{Q}}_{hvac}, \dot{\mathbf{V}}] \leftarrow \text{Optimize}(\mathbf{State}(i), K)$
 $[\dot{Q}_{opt}, \dot{V}_{opt}] \leftarrow [\dot{\mathbf{Q}}_{hvac}(1), \dot{\mathbf{V}}(1)]$
 $[\mathbf{C}_r(i+1), \mathbf{T}_r(i+1)] \leftarrow \text{Simulate}(\mathbf{State}(i), \dot{Q}_{opt}, \dot{V}_{opt})$
end

At the beginning of this MPC algorithm (Algo. 1), room CO₂ concentration and temperature vectors $\mathbf{C}_r, \mathbf{T}_r$ are initialized with the current room conditions c_r, t_r . Then the

algorithm runs for a total of n steps. In each step i , the optimization function considers a K -step outlook into the future and returns the HVAC Energy Flow $\dot{\mathbf{Q}}_{\text{hvac}}$ and Volumetric Flow Rate $\dot{\mathbf{V}}$ which are both K -dimensional vectors corresponding to the overall optimal decisions in the next K steps subject to the given objective function and constraints. Then the first decision vector $[\dot{\mathbf{Q}}_{\text{hvac}}(1), \dot{\mathbf{V}}(1)]$ is considered as the decision made at time i and is used to simulate the room CO₂ concentration and temperature for the $(i + 1)$ -th time step, i.e., the next time step.

In the meantime, this process involves a large number of decision variables, which, if formulated inefficiently, could consume huge computation power. As room temperature dynamics and CO₂ mass transfer are highly coupled, we introduce the following techniques to convert this non-convex problem to a convex system. Then we use *CVX* [42, 41], a convex problem solver, to optimize MPC control decision.

5.3.1 Convex Relaxation: Discretize Differential Equations

Equation (5.4), the discretized analytical solution to Equation (5.3) is highly non-convex, as it involves multiplication of two variables: \dot{V} and b that depends on \dot{V} . Therefore, we directly discretize Equation (5.3) which gives an approximated change of CO₂ concentration with respect to a small time interval Δt

$$\Delta C_r(k) = \left(\lambda(k)C_{out} - \lambda(k)C_r(k) + \frac{\dot{E}(k)}{V} \right) \Delta t \quad (5.19)$$

The approximated C_r then can be calculated as:

$$C_r(k + 1) = C_r(k) + \Delta C_r(k) \quad (5.20)$$

5.3.2 Convex Relaxation: Small Time Interval Approximation

Note that in a simulation of K time steps, Equation (5.19) still introduces non-convexity, as the terms C_r , which depends on the previous time step control decision, and λ , are still connected. Therefore, we change the dependent variable $C_r(k)$ to a parameter by assigning it a constant value $C_r(1)$, which is obtained from the previous simulation step. Conceptually, we assume that the CO₂ concentration of air flowing out of the room only depends on the initial (last time step simulated) concentration, and ignore the dynamical behavior in the next K time steps. This gives

$$C_r(k + 1) = C_r(k) + \left(\lambda(k)C_{out} - \lambda(k)C_r(1) + \frac{\dot{E}(k)}{V} \right) \Delta t \quad (5.21)$$

We use the same method when encountering the coupling in Equations (5.2) and (5.6).

We define the dependent variable $T_r(k)$ in Equation (5.6) as a constant, equal to the target temperature T_t .

$$\dot{Q}_{out}(k) = \rho_{air} \dot{V}(k) C_{p,air} (T_{out}(k) - T_t) \quad (5.22)$$

When the room is occupied, room temperature is optimized to reach target temperature deadband $T_t - 0.5 \leq T_r \leq T_t + 0.5$ with thermal comfort penalty. $\dot{Q}_{out}(k)$ is mostly driven by change in $\dot{V}(k)$. When room is unoccupied, on the other hand, $\dot{V}(k)$ almost follows \dot{V}_{min} , which renders $\dot{Q}_{out}(k)$ negligible.

5.3.3 Convex Relaxation: Linear Approximation

Equation (5.9) is linearized using linear approximation which gives

$$H(k) = f(k) \quad (5.23)$$


for $f(k) \in [0, 1]$.

5.4 Results and Discussion


5.4.1 Convex vs. Non-Convex Simulation Comparison

To validate the convex relaxation process, we use Algo.1 which optimizes with convex simulation and simulates with non-convex systems. Results in Figure 5-1 demonstrate that, when the optimization horizon k is set to 24 steps with each step being 5-minute (2 hours ahead), both simulation methods yield comparable results in all variables of interest. This is because the optimization and simulation takes place alternately every 5 minutes; any inaccuracy or deviation incurred by convex approximation can be corrected in a timely manner by running a round of non-convex simulation.


However, these convex relaxation methods restrict us from looking further into the future, i.e. having a larger optimization horizon k , as some of them only guarantee local approximation. With a larger horizon, we could possibly get an infeasible solution as some of the constraints no longer hold. A stable and well-defined convex relaxation algorithm, such as Taylor Series expansion, can address this issue.




(a) Predicted (convex) vs. simulated (non-convex) room temperature



(b) Predicted (convex) vs. simulated (non-convex) room CO₂




(c) Predicted (convex) vs. simulated (non-convex) DOAS power




(d) Predicted (convex) vs. simulated (non-convex) CO₂ emission


Figure 5-1: Convex vs. non-convex simulation result comparisons



(a) Power consumption vs. grid carbon intensity



(b) Room temperature vs. grid carbon intensity



(c) Room CO₂ level vs. grid carbon intensity

Figure 5-2: MPC response to variation in grid carbon emission intensity²

²All plots begin at 7:00 am; the room starts to get occupied at 9 am (540 min) and becomes unoccupied after 6 pm (1080 min); during the room operating hours, room occupancy rate varies according to class schedule. The dash lines in (b) and (c) refer to the deadbands between which no penalty incurs.

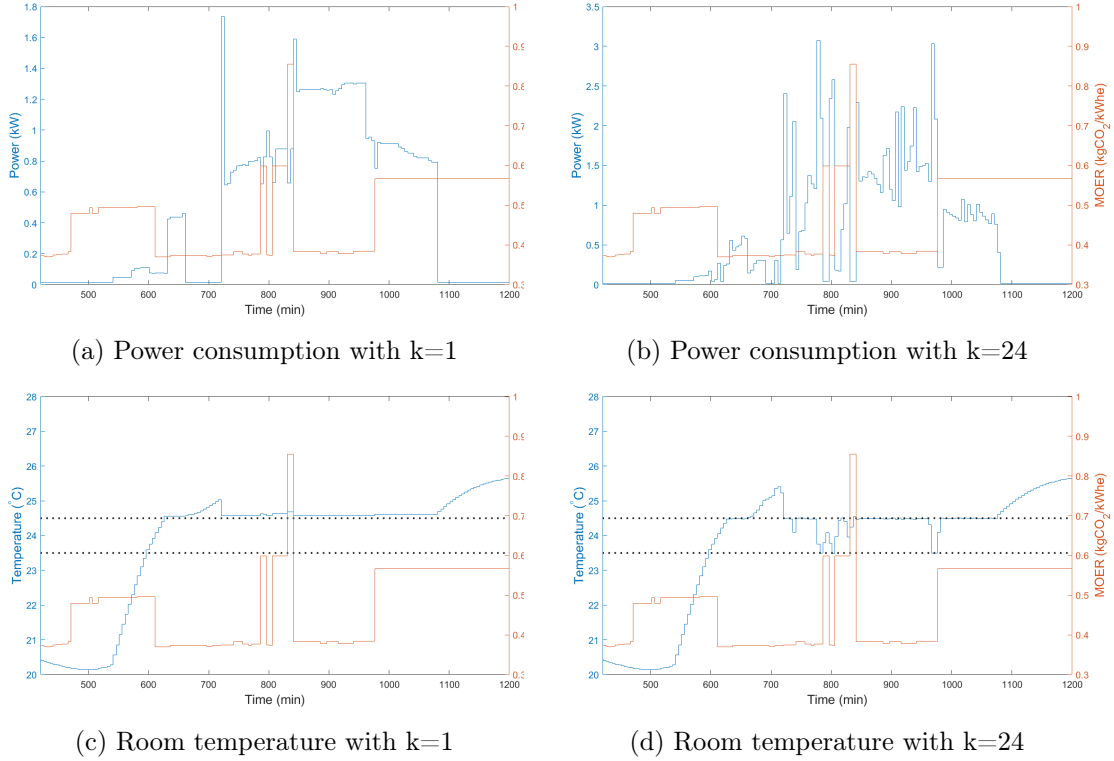


Figure 5-3: Control decision with different optimization horizons

5.4.2 Model Response

With a 2-hour optimization horizon, MPC makes predictive control decisions in response to forecast changes in the electricity market and grid operation. Figure 5-2 shows the model response to variation in grid carbon emission intensity. The right y axis in each figure denotes MOER. The two horizontal dot lines in Figure 5-2b and Figure 5-2c correspond to the thermal comfort and IAQ range respectively, within which the system will not be penalized. When the algorithm foresees a sudden spike in grid MOER (from 800 to 850 mins in simulation horizon), it plans to reduce power consumption in this period. To achieve this, it decides to pre-cool the room so that room temperature and CO_2 level stay close to and even go below their corresponding lower bounds when the grid MOER is relatively low. This ensures almost zero power consumption during MOER peak period. Consequently, room temperature and CO_2 level drift up to and could be even above upper bound when MOER peaks. Depending on the weight assigned to each optimization objective, the algorithm decides if it takes priority in minimizing total carbon emission or satisfying thermal and IAQ constraints. Note that periods during which room temperature drifts up and CO_2 drops down (from 650 to 700 min) are unoccupied in working hours. Thus, no penalty is imposed on room temperature and CO_2 level.

5.4.3 Greedy Optimization vs. Optimization with Foresight

As described in the previous section, based on the availability of knowledge about the future, MPC plans an optimal control schedule in a different manner. Take two extreme cases as examples shown in Figure 5-3. For easy visualization, all figures are cropped to only display the time interval between 8:30 am (500 min) to 8:00 pm (1200 min). The left two figures 5-3a and 5-3c demonstrate the optimized power consumption and simulated room temperature with a horizon of $k = 1$ (5-min ahead). The right two figures 5-3b and 5-3d showcase results with a horizon of $k = 24$ (2-hour ahead). Other model parameters are kept the same.

In both cases, the algorithm attempts to minimize power consumption during peak MOER period. In the case of $k = 1$, power consumption only goes down by 0.2 kW and immediately increases to 1.6 kW after the MOER peak. In the case of $k=24$, when foreseeing the sudden rise in MOER, the algorithm is able to cut off power consumption almost perfectly. On the room temperature side, temperature is always maintained above thermal comfort range with one-step ahead planning, whereas temperature with 2-hour planning mostly falls into the thermal comfort band. In addition, 2-hour ahead planning shows a more obvious pre-cooling strategy which gives the algorithm additional capacity of dealing with MOER spikes without incurring large penalty on thermal comfort. Admittedly, in the 2-hour prediction scenario, the control decision involves more fluctuations and therefore is more sensitive to changes in the system. To overcome this phenomenon, constraints that control the frequency and magnitude of changes in the control decision can be imposed in the system.

In theory, a larger horizon gives MPC more capacity in achieving load-shifting and GHG and cost reduction goals. However, in reality, it might not be feasible to get an accurate time-series profile forecast of certain parameters in the system. For example, except room schedule, weather forecast, and 24-hour ahead LMPs, grid MOER forecast is only available for two hours ahead, which limits the maximum prediction horizon.

5.4.4 Multi-objective Pareto Front

The MPC optimization involves four objectives: (1) total carbon emission, (2) total energy cost, (3) thermal comfort, and (4) indoor air quality, of which the latter two terms are formulated as penalty functions. We use a weighted sum method to obtain the Pareto Front in the design space. Specifically, the four objectives are grouped into two subgroups, with (1) and (2) being the Cost group (both economically and environmentally) and (3) and (4) being the Performance group. Each objective is scaled to the same order of magnitude beforehand so that the algorithm is able to react to the changes in weights. As is described in Algo.2, a set of weights, w_1 and w_2 , is assigned to the Cost group and Performance group

respectively³. We use the same MPC algorithm with weights assigned to each objective, which outputs total energy cost Tot_{Pr} , total carbon emission Tot_{Emit} , cumulative thermal comfort penalty Tot_{Pt} , and cumulative IAQ penalty Tot_{PCO_2} .

This Pareto Front analysis is performed on the horizon of $k = 1$, as shown in Figure 5-4. It is evident that there exist trade-offs among four objectives. In general, as the penalties are decreased with a shift of focus on optimizing performance, total carbon emission and energy cost increase. Such a trade-off has been the subject of discussion among building engineers and health scientists. On the one hand, building engineers and environmentalists aim to design and develop more energy efficient and sustainable buildings; on the other hand, building owners and dwellers would like to live in a healthy and comfortable environment. Investigating the Pareto Front allows us to quantitatively understand this trade-off.

Specifically, in this example analysis, the impacts of improvement in IAQ and thermal comfort on room operation cost are different. As depicted in Figures 5-4a and 5-4b, a 70% improvement in IAQ Penalty only introduces 10% increase in both carbon emission and energy cost. On the other hand, a reduction of thermal comfort penalty by 50% incurs approximately 35% increase in both. Note that the result is specific to buildings in Massachusetts only on a selected day, and model parameters (outdoor weather, occupancy schedule, grid MOER and so on) vary hourly, daily and seasonally. If a more general conclusion is needed, a whole-year analysis is required to check the impact of building performance on operation cost. In addition, instead of lumping four objectives into two groups, more advanced techniques can be used to explore the design space by assigning each objective with an individual weight.


Algorithm 2: WEIGHTED AVERAGE PARETO FRONT

```


w1 ← [0.01:0.01:0.09, 0.1:0.05:0.95]
w2 ← 1 − w1
for  $i \leftarrow 1$  to  $len(\mathbf{w}_1)$  do
    [Pr, Ec, Pt, PCO2] ← MPC( $[c_r, t_r]$ ,  $K$ , w1( $i$ ), w2( $i$ ))
    [ $Tot_{Pr}$ ,  $Tot_{Emit}$ ,  $Tot_{Pt}$ ,  $Tot_{PCO_2}$ ] ← [ $\sum \mathbf{Pr}$ ,  $\sum \mathbf{E}_c$ ,  $\sum \mathbf{P}_t$ ,  $\sum \mathbf{P}_{CO_2}$ ]
end

```

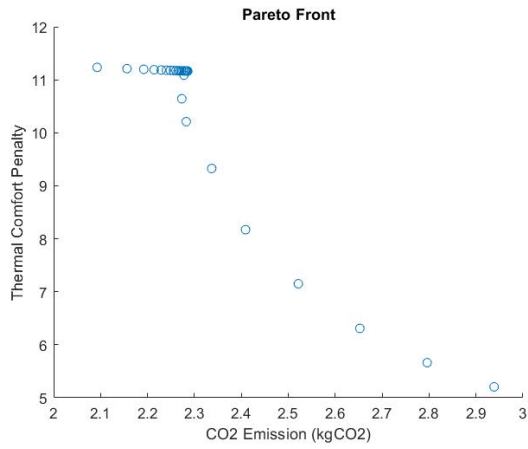
³Note that weights are not sampled at equal distance as each objective group has different sensitivity to changes in weights.




(a) IAQ penalty vs. carbon emission



(b) IAQ penalty vs. power consumption



(c) Thermal comfort penalty vs. carbon emission



(d) Thermal comfort penalty vs. power consumption

Figure 5-4: Pareto front of multi-objective optimization




Figure 5-5: Model robustness⁴

⁴All plots begin at 0 am; the room is scheduled to become occupied at 9 am (540 min) and become unoccupied after 6 pm (1080 min); from 11 am to 12 pm, room is scheduled to be unoccupied.

5.4.5 Final Design Solution

As discussed in the Multi-objective Pareto Front section, the Pareto Front analysis on a specific experiment day cannot generalize the true interactions among objectives. Connections and interactions among objectives cannot be determined without running a whole-year simulation. In addition, multiple factors could influence the final design choice: (1) building owners and users: whether building owners value occupant comfort more than operational cost and environmental impact; (2) building types: whether the type of activities conducted in the building or building function has specific requirement on room conditions; for example, laboratories may have strict temperature control; (3) number of future horizons to look into: whether sufficient and accurate predictions over future model parameters are readily available.

5.4.6 Model Robustness

When running this MPC model, we use historical data as predicted model parameters, which provides 100% knowledge into the future as a fundamental modeling assumption. Therefore, it is necessary to examine model robustness when inaccurate information is given to the system, which could introduce discrepancies between predicted and actual observed room conditions. So we introduce occupants in the scheduled-unoccupied period (1 hour before 9 am, 11 am to 12 pm, and 30 mins after 6:30 pm). Because the algorithm is unaware of the real room occupancy status, it continues to optimize with constraints and parameters designed for the unoccupied period. As is shown in Figure 5-5, room temperature increases with unexpected increase in occupancy; however, in such a short time interval, room temperature does not deviate too much from desired thermal comfort band except in the morning when occupants might feel too cold. Among all uncertainties involved in future scenario prediction, tracking stochastic occupant activity may be considered as the most difficult task. To mitigate such stochastic effects on model performance, motion sensors can be installed to track real-time occupancy density to inform the algorithm.

Chapter 6

Experiment Setup and Design

To validate the proposed control strategies developed in simulation and optimization models, we implement the two-step optimization method in two MIT classrooms as onsite experiments where we control the room setpoints and observe the resulting load conditions. The onsite experiment is deemed critical because we realize that some previous research either report optimization results in simulation settings, or rely on simulation software generating baseline load as the ground truth which usually involves an uncalibrated performance gap. The goal of this series of experiments is to control real building operations and draw generalizable conclusions for the two-step optimization approach.

6.1 Test Room Configuration

The onsite experiments take place in Rooms 154 and 160, the two classrooms pre-selected by the research team, in Building 66, the chemical engineering building on campus. As is shown in Figures 6-1 and 6-3, two rooms are located on the first floor, with large windows facing southeast. Classes are held during weekdays in the semester according to course schedules. The maximum occupancy capacity of both rooms is approximately 40 students. The two rooms are almost identical in terms of the room size and room HVAC equipment setup.

6.1.1 HVAC Equipment Setup and System Configuration

The two spaces have identical HVAC system configurations, both equipped with a combined system of Variable Air Volume (VAV) with reheat, chilled beam, and fin tube radiator (FTR), providing two-tiered conditioning. In this building, VAV supplies 100% outdoor air, pre-conditioned by the central air handling unit (AHU). As is shown in Figure 6-2, the supply air volume is modulated by a damper. The air goes through the heat exchanger controlled by a reheat valve. Then the air is sent to the chilled beam system where heat exchange takes place between chilled water and air, controlled by a cooling valve. There is



Figure 6-1: Building 66 floor plan

a separate FTR to meet the room heating load.

The coordination between the chilled beam and VAV follows a two-step procedure. In the cooling mode, the cooling load is first fulfilled by the chilled beam¹. After the chilled beam reaches the maximum capacity, the additional cooling load is met by the VAV box. In the heating mode, depending on the ambient temperature, the VAV reheat and FTR may work sequentially or in unison. The detailed control logic is included in Appendix C.

Figure 6-2 shows a Building Management System (Ecostructure) control dashboard developed by Schneider Electric. The dashboard monitors real time system status, such as valve positions, supply air volume, and air temperature, and allows for inputting heating and cooling setpoint temperatures. In addition, occupancy sensors are installed in the two rooms, providing researchers with detailed information about real-time room occupancy status and helping validate experiment results.

6.2 Experiment Design and Methodologies

6.2.1 Control and Baseline Test Sequence Design

To accurately assess the GHG emission savings through building HVAC operations, one would ideally need a parallel universe test, with one test implementing the optimal control solution and the other operating as business-as-usual. While there's no perfect parallel universe in reality, we have considered four candidate solutions that emulate this effect.

¹Note that the maximum cooling load provided by the chilled beam is approximately 30% of the design capacity to prevent condensation. This is because the existing chilled beam is not equipped with drainage.

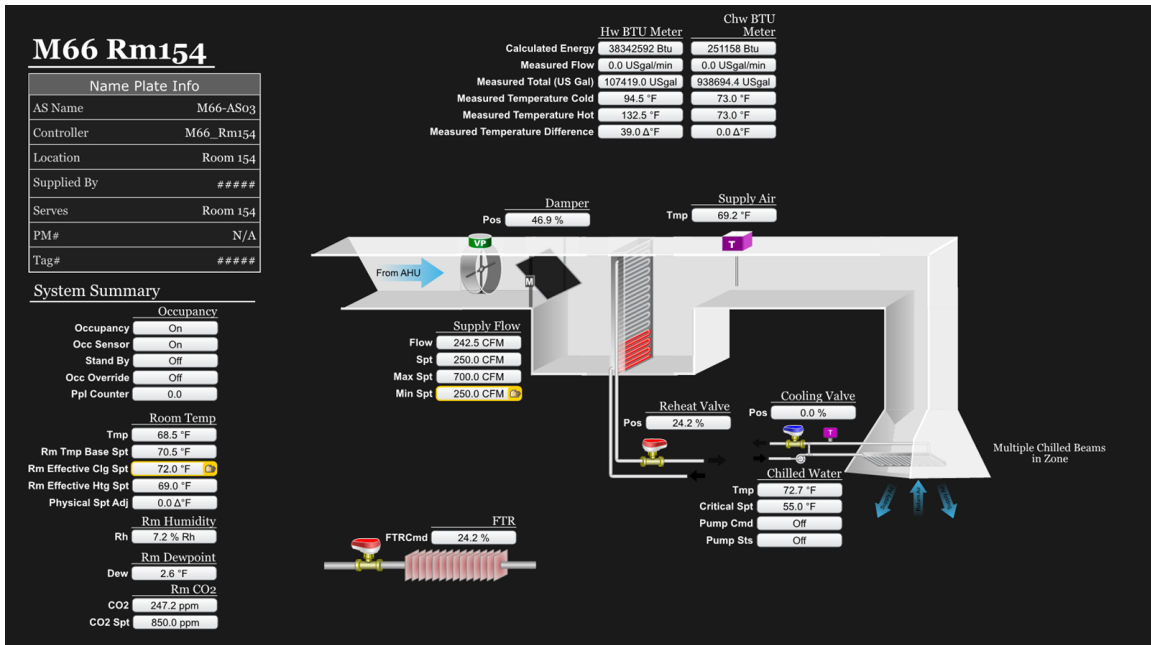


Figure 6-2: Room 154 BMS control dashboard

Approach 1: Sequential Test

The first approach is to design a back-to-back experiment in one room that puts one baseline and one optimal control test in sequence. For example, a two-day experiment during which the first day runs as the baseline and the second day as the optimal control could serve this purpose. Since the two experiments take place within the same space, it is certain that the room thermal properties and thermal response are consistent. However, this approach is more susceptible to variations of external conditions, such as the ambient temperature, solar radiation, and occupant behaviors, that would introduce different load conditions and make the results less comparable.

Approach 2: Parallel Test

The second option is to design a parallel test in two identical rooms. For example, a baseline and a control test are conducted in two rooms simultaneously. In this case, Room 154 is chosen as the experiment room, while Room 160 is kept as the baseline. The parallel test requires the two rooms to have a high degree of similarity in room size, orientation, location in the building, occupancy condition considering occupant loads, and HVAC equipment type and control sequence, so that the two rooms would have identical thermal response if the same control sequence was sent to the two controllers. To use the second approach, we would need to first verify the similarity of thermal response of the twin classrooms, which we will refer to as “baseline verification test.” The following section discusses the details of a verification test and Chapter 8 concludes the validity of the test results.

Approach 3: Scaled Predicted Baseline

The third option involves fewer restrictions on room conditions and system setups, as well as external environment. Instead of using Room 154 or Room 160 as a direct baseline reference, this approach investigates the consistency of the relative load changes in Room 160 and Room 154 within themselves, and adjusts the difference in Room 154 energy use (optimal control vs. baseline), by the percentage difference in Room 160 energy use over the two days. For example, if we observe the two rooms follow the same load trends, and a 10% load change happens between the first and second baseline tests in Room 160, we would scale the load of Room 154 by the same percentage as the baseline for the second day and compare it against the optimal control load to estimate load savings. To verify the similarity in load trend, it is critical to identify a period during which room conditions such as occupancy and lighting remain similar to avoid disturbance. The comparison result is presented and analyzed in Chapter 8.

Approach 4: Machine Learning Driven Baseline

The last option is to establish a virtual baseline through a predictive algorithm such as a machine learning model. Given the room conditions and external environment will not always be the same, which may lead to disturbances in baseline reference generation, an accurate load prediction model could fix the gap. This approach is especially useful when the optimal control method is scaled to the entire campus where parallel or sequential baseline tests are no longer a viable option. However, the challenge of adopting this approach is its high demand in prediction accuracy. The error or noise introduced from the baseline prediction should be negligible compared to the savings, to allow for the following saving estimation.

6.2.2 Experiment Protocol

Except for the first and last experiments conducted during the summer break and winter break when campus are closed, the rest of experiments take place over multiple weekends during the semester. To perform controlled experiments and exclude the effect of confounding variables on the experiment results, such as occupant behaviors, we purposefully choose weekends when occupant activity is less significant.

Over several weekend and holiday experiments, we observe room load trends with varying degrees of control of occupant behaviors². For some weekends, students are free to group and self-study in the rooms without being aware of the experiment. On other weekends, we reserve the space through the Registrar Office and put signs (attached in Appendix A) prohibiting students entering rooms, opening connecting doors, changing blind positions, or

²The detailed experiment protocol and setup is shared in the Appendix A.

turning on projectors. Recognizing students’ need for spacious study space especially during the pandemic, we create room usage instructions that request students to restore experiment conditions after their study. We also setup illuminance sensors to measure room lighting conditions and quantify the difference in solar heat gains between the two rooms.

Each set of experiments lasts for two to three days and follows the schedules below. The majority two-day tests follow the experiment sequence in Table 6.1a. The three-day tests follow the schedule in Table 6.1b.

Table 6.1: Baseline and optimal control test schedule for a two-day test period (a) and a three-day test period (b)

(a)

	1st experiment day	2nd experiment day
Room 154	Baseline	Optimal control
Room 160	Baseline	Baseline

(b)

	1st experiment day	2nd experiment day	3rd experiment day
Room 154	Baseline	Optimal control	Rule-based
Room 160	Baseline	Baseline	Rule-based

The test schedules outlined in Tables 6.1 serve multiple purposes and can be applied to different control and baseline test approaches described in Section 6.2.1. To validate the sequential mode, we compare the baseline and optimal control loads in Room 154 only. To validate in parallel mode, the first day is used for validating room baseline load by comparing thermal loads for Rooms 154 and 160. In the next day’s cross-room comparison, Room 160 serves as the baseline to Room 154, which allows for a direct comparison between the two load profiles and calculation of the savings. To validate the scaled predictive approach, multiple baseline tests in Room 160 and Room 154 are compared against each other to define the load relationship. Based on the relationship, Room 154’s baseline on the second experiment day can be predicted with information of Room 160’s load and compared with the observed optimal control load.

The only difference between a two-day test and a three-day test is that the latter includes a rule-based control test day. The algorithm outputs temperature setpoints only based on the room occupancy schedules, which is a common practice in typical building operations. In the following sections, we refer to the first-day experiment as the “twin test room validation”, and the second day as the “optimal setpoint implementation”.

6.2.3 Data Requirement

In preparing and conducting experiments, the following data listed in Table 6.2 are collected for analyzing heating and cooling loads in Rooms 154 and 160 @Building 66. APIs are used



Figure 6-3: Classroom configuration

to retrieve data from various databases. Note that we don't have electricity consumption data on the room level.

Table 6.2: Data types and databases for experiment setup and result analysis

Data type	Database	Notes
BTU meter	AI Report Server (now migrated to Ecostruxture)	Heating and cooling loads from chilled beams and FTR
VAV energy	Clockworks	VAV related data, such as VAV supply and return air temperatures, air flowrate, and outdoor temperature, etc.
Occupancy	Density with API	Room traffic
Grid carbon intensity	ElectricityMap with API	Real time and forecasted grid carbon intensity
Weather	Oikolab with API	Historical and forecasted weather, including parameters such as dry-bulb temperature, solar radiation, and relative humidity, etc.
Room setpoint control	Ecostruxture	Built-in dashboard to define 24-hour setpoints for experiments
Thermal comfort feedback	AWS	Feedback from occupants regarding indoor thermal comfort, to be integrated with future experiments

6.2.4 Experiment One: Twin Test Room Validation

To validate that the two rooms have similar thermal response and room loads are comparable, which is a major hypothesis in adopting a parallel baseline, multiple set of experiments are conducted, with different room control conditions. in cooling, heating and shoulder seasons. In most experiments, the cooling and heating setpoints are set at 72 °F (± 0.5 °F) and 69

°F (± 0.5 °F), respectively, in both rooms.

As is illustrated in Figure 6-3, factors that may influence the classroom loads are: (1) lighting condition, (2) external conditions, such as shading from trees and buildings nearby, (3) front and side door³ openings, (4) blind shading condition, and (5) projector on/off status. Specifically, two side door openings can cause unintended heat transfer between the adjacent rooms, and the same applies to the front door that exchanges room air with the hallway. In addition, both rooms have one translucent and one opaque blinds to reduce solar glaring. Occupants can lower the blind to a specific position (i.e. 0%, 25%, 50%, 75%, and 100%) at any time at their will to reduce direct solar radiation or glaring, and have more room privacy.

6.2.5 Experiment Two: Optimal Setpoint Implementation

In the second day of the experiment, Room 154 implements the optimal setpoint schedule, while Room 160 keeps the constant cooling and heating setpoints at 72 °F (± 0.5 °F) and 69 °F (± 0.5 °F) respectively.

The thermal load of the room (Q_t) is split into three portions: load (1) provided by the VAV supply air (Q_{VAV}), (2) chilled beams (Q_{CB}), and (3) FTR (Q_{FTR})⁴.

$$Q_t = |Q_{VAV}| + |Q_{CB}| + |Q_{FTR}| \quad (6.1)$$

The VAV supply load can be estimated by:

$$Q_{VAV} = c_{air} \dot{m}_{air} \Delta T = c_{air} \dot{m}_{air} (T_{room} - T_{supply}) \quad (6.2)$$

Where m_{air} is the air mass flowrate, T_{room} the room temperature measured by room thermostats, equivalent to the return temperature assuming well-mixed condition, T_{supply} is the supply air temperature before entering the chilled beam. These parameters are directly measured by sensors installed inside the rooms or ducts. c_{air} is the specific heat capacity of air under constant pressure condition, which is 0.240 BTU/°F-lb (1.012 J/g-K). Air density at 60 °F and 1 atm is 0.0754 lb/cu-ft (1.208 kg/m³).

The cooling load data from chilled beams and heating load data from FTR are directly extracted from BTU meter database, which is estimated by calculating the heat extracted by water through heat exchange between air and water, assuming 100% heat exchange efficiency:

$$Q_{CB} = c_w \dot{m}_w \Delta T = c_w \dot{m}_w (T_{ws} - T_{wr}) \quad (6.3)$$

where c_w is the heat capacity of water, \dot{m}_w is the water mass flowrate, T_{ws} is the supply

³There are three doors in each room, one front door, and two connecting doors to adjacent rooms.

⁴The simultaneous heating and cooling (positive and negative) loads from different systems should be added together, instead of cancelled out.

water temperature in the pipe, and T_{wr} is the return water temperature.

6.2.6 Experiment Constraints

The differences in the actual experiment setting and the simulation require us to fine tune our implementation strategies to ensure the results are comparable and interpretable. The two noticeable factors that need to be addressed before implementation are (1) the setpoint control logic and (2) the ability to track setpoint in the instrumental stage.

Dual vs. Single Setpoint Temperature Control

The test rooms adopt a dual-setpoint control that defines a lower and upper bound of room temperature. With this control scheme, both heating and cooling setpoints can be active in any season and any time of the day. The two setpoints form a deadband between which the heating or cooling system remains off. This setup is different from our simulation model where the room is controlled by a single setpoint. To illustrate, in cooling seasons, regardless of how far the room temperature drifts below the cooling setpoint, heating will never be triggered with a single setpoint in the simulation model. The drifting period allows the room to be cooled down without incurring heating energy consumption and therefore achieves load savings while pre-conditioning the space for the following day.

However, in the case of Rooms 154 and 160 with dual setpoints, both heating and cooling equipment will be triggered anytime the room temperature does not lie in the deadband. This has two impacts on our experiments: (1) the drifting period simulated and optimized by the algorithm could not be reflected in the experiment as the dual setpoint will prohibit temperature drift outside the deadband, and (2) what is worse, the purposeful drift combined with dual setpoints will incur a higher heating and cooling energy consumption.

Ability to Track Setpoint Schedule

The simulation and optimization model relies on the hypothesis that the room load demand can be fulfilled immediately with HVAC equipment. This has two implications. First, the room temperature will closely follow the designed setpoint schedule. Second, the delay in system response can be fairly neglected even when setpoint perturbations are relatively large, i.e., 5 °F⁵.

However, such a hypothesis is not perfectly compatible with the HVAC operation in the two rooms. On the one hand, the dual temperature setpoints create a control deadband that allows a bigger shift in the room temperature as long as it stays within the range. The larger the deadband, the less accurate control over room temperature we have. On the other

⁵Recall the two-step approach involves a 5 °F setpoint perturbation in simulation.

hand, through the initial system commissioning, we identify an unexpected longer system response time for the cooling system that significantly violates our modeling assumptions.

6.2.7 Two Optimization Framework and Targets

Given the limitations on room-level equipment control and modeling differences in optimization framework, we have proposed two approaches to estimate GHG emission savings, first from calculated room thermal load, and second from mid-stream load consumption from heating and cooling energy streams. Figures 6-4 and 6-7 illustrate two setpoint control strategies that aim to minimize their associated energy and load.

Approach 1: Optimize for Thermal Load

This approach only considers room load based on the resulting room condition and supply air temperature from the three systems, and disregards the intermediate conditioning procedure that may involve more energy consumption. With this approach, we assume the room load serves as a proper reference for energy consumption, which is similar to our modeling approach via EnergyPlus. More specifically, the proposed framework assumes the room temperature closely follows the setpoint schedule, and it does not account for any reheating process caused by the dual setpoints in practice.

To apply the first approach, we narrow the setpoint deadband to the minimally allowable gap. This renders the dual setpoints to act as a single setpoint. To prevent short-cycling, we choose to work with a 0.5 - 1 °F gap between cooling and heating setpoints in the experiment room. In doing so, if everything operates correctly, the room temperature should fall within a 0.5 - 1 °F band, shown in Figure 6-4, which fulfills our first modeling assumption.




Figure 6-4: [Room 154] Cooling, heating setpoints and actual room temperature for 24-hour horizon on Sep. 23. The cooling (blue) and heating (red) setpoints are set at 70 °F and 69 °F respectively, and the actual room temperature is in grey.

Figure 6-5 illustrates the adverse effect of increased energy consumption, associated with a narrow deadband. During the experiment 24-hour period (in a cooling season), in addition to the cooling system, the heating system is triggered on to ensure room temperature is maintained between 69 °F and 70 °F,. More specifically, the reheat valve for VAV is open for approximately 50% time to heat up supply air to above 55 °F, the supply air temperature from the AHU to provide heating. We can conclude that using a narrower deadband with a dual setpoint control to simulate a single setpoint control produces counter-effect on experiment results. In fact, the narrower the band, the larger the energy consumption from both heating and cooling load would become. With such a framework, minimizing energy with perfect tracking of room temperature is inherently incompatible.




Figure 6-5: [Room 154] The valve positions of VAV reheat coil and chilled beam on Sep. 23

Approach 2: Optimize for Mid-stream Load Consumption

The second approach considers the mid-stream load consumption, that is, accounting for both heating and cooling loads used for maintaining room temperature. Due to the dual setpoint control, the reheat system will be triggered when room temperature is below the heating setpoint even in the cooling season. While a constant 3 °F deadband in the existing control could help reduce energy usage by creating a large band, there's further opportunity to fully utilize the benefits of dual setpoint control – to deploy a varying setpoint deadband.

In a constant deadband control, heating and cooling setpoints move simultaneously, as is shown in Figure 6-6 in a heating season experiment. It is obvious that from midnight to early morning, and evening to nighttime, the optimization algorithm intends to lower heating setpoints to reduce heating load in a heating season. However, the constant deadband also forces a lower cooling setpoint, which in turn demands extra load to make space colder, even if the intention is to make the temperature float without incurring heating energy.

To overcome the unintended outcome, a slight modification to the dual setpoint creates




Figure 6-6: [Room 154] Heating and cooling setpoints on Oct. 24 in the heating season

more potential for energy savings. As is illustrated in Figure 6-7, if we fix the cooling setpoint and only adjust and optimize for the heating setpoint, this will avoid over-cooling the room and allow for larger temperature fluctuations. The cooling setpoint can be chosen based on survey results of occupants' thermal comfort, designating the highest allowable room temperature, in case heat injection from occupants and solar heat gains render the space unsuitably warm. The heating setpoint, under most circumstances, will regulate the room temperature in a heating season.




Figure 6-7: Hypothetical dual setpoint control with a varying setpoint deadband, an example for heating season

Recall that the proposed optimization algorithm relies on the assumption of room temperature closely tracking setpoints. This requirement can be fulfilled in cooling and heating seasons because cooling and heating setpoints often dominate the temperature control in the respective seasons even within a dual setpoint control system. For example, in the heating season, room temperature most of the time follows heating setpoints, and cooling loads are

not prevalent. However, we should be careful that this varying setpoint control method may not satisfy such an assumption during the shoulder season when heating and cooling interchanges frequently throughout a day. Additional attention needs to be taken to decide which setpoint would be active during a specific period, to increase the flexibility for room temperature drift in the right direction.

Chapter 7

HVAC System Commissioning

To ensure the experiment results are reliable and replicable, commissioning of HVAC systems, hardware, and software was conducted in heating and cooling seasons. The purposes of the tests include:

1. To understand the control sequences of three systems collaboratively in supplying heating and cooling loads
2. To test the system's response time and ability to track setpoint temperatures
3. To identify any other potential issues that may prohibit the completion of future experiments

The following sections present a detailed description of the hardware and software issues; some were fixed immediately after being identified, while others are campus-wide problems that need to be discussed further to better serve future experimental needs and improve building operation efficiency.

7.1 Inability to Maintain the Setpoints

Failure to track room setpoints has been observed throughout multiple test periods. An example of an experiment on Sep. 8 in a cooling season is shown in Figure 7-1. While room temperature is kept within the heating and cooling setpoints during most of nighttime, a significant deviation occurs from ~9:30am to 6pm, highlighted in red, with the room temperature peaks at 76 °F in the early afternoon. This shows the limited cooling capacity of the VAV and chilled beam systems, unable to meet the increased cooling load during the daytime. Noticeably, such condition also happens in the adjacent Room 156. While this challenge only occurs in the cooling season, it needs to be resolved urgently because without the ability to track setpoints, the system would not be able to follow the designed optimal

setpoint schedule which is our main lever to save GHG emission and energy consumption, and therefore may not achieve the desired research outcome.




Figure 7-1: [Room 160] Setpoints and room temperature for 24-hour horizon on Sep. 8

7.2 Constant Supply Air Volume (CAV) Operations in Variable Air Volume (VAV) Systems

After further investigation, we identified that the system’s failure to fulfill the required cooling load partially comes from the altered operation of a variable air volume (VAV) system. While the system was initially designed to provide variable air flow to meet different load demand, it in practice operates as a Constant Air Volume (CAV) system where the supply air volume is fixed at 450 cfm (cubic foot per min).

From our conversation with Schneider Electric, the supply air flow was initially capped at 900 cfm and floored at 450 cfm, but they were requested to reduce the maximum air flow to 450 cfm because occupants reported noise of equipment operation. As a result, the same upper and lower bound of air flow render the system running as a CAV. It was revealed from the engineering team that the 450-cfm minimal outside air flow design came from the requirement of indoor air quality standard, or to maintain a healthy room CO₂ level.

Now that VAV supply air mode has turned into a CAV operation, the flexibility of heating and cooling load provided by supply air relies only on adjustment of the air temperature. For example, to maximize cooling supply during the peak hours in summer, the supply air is kept at ~55 °F, the lowest possible supply air temperature provided by the air handling unit (AHU). During the nighttime, when the cooling load is smaller, the supply air gets reheated by a VAV reheat coil at room level. By changing the supply air temperature, rather than supply air volume, the system keeps the room condition within a designed temperature range.

Note that this converting the system to a CAV took place before COVID-19, and is not a temporary control override in response to COVID-19. However, this is neither an effective nor energy efficient control logic. A typical VAV system would supply air at constant temperature most time and vary supply flowrate to meet different room loads. For example, in a cooling season, a higher cooling load incurs a larger flowrate. Since the system virtually operates as a CAV unit, its only capacity in adjusting loads is changing the supply air temperature through the reheat coil. On the one hand, this largely limits its cooling capacity which is defined by the lowest supply air temperature and the fixed airflow. On the other hand, it relies on reheat coil to heat up supply air when room cooling load is small, which consumes additional energy. Hence, we seek to relax the flowrate constraint by extending the minimum to 250 cfm and maximum to 700 cfm. In this way, the system gains back its original variable flow feature, increases its capacity in adjusting room temperature, and improves the overall energy efficiency. Nonetheless, dropping the minimal outdoor air flow may not meet the ASHRAE 62.1 [18] indoor air quality standard. Based on the design standard, 450 cfm outdoor air would satisfy a class size of 36 students¹ and a room size of 760 ft². Lowering the flowrate to 250 cfm limits the system's capability to only accommodate 16 students. We understand the temporary control change is not ideal and recommend future experiments holistically revisit the system design and come up with a well-rounded plan that satisfy both experimental and practical constraints.

While this CAV issue has been resolved in the two experiment rooms, other rooms in Building 66 and other campus buildings may face similar issues. It is not a surprise that other campus buildings may fail to meet setpoint tracking requirements. Indeed, to ensure the participation of most campus buildings in the setpoint adjustment, this is the first and foremost task the team should address. Fundamentally, it is critical to understand (1) what other buildings on campus have similar issues, (2) what happened historically that has led to the changes in building operations, (3) how feasible it is to address such a campus-wide issue.

7.3 Limited Cooling Capacity of Chilled Beams

Although chilled beams have been proved to provide better thermal comfort to the occupants [67], they have a slow response rate and small cooling capacity due to design constraints. As is shown in Figure 7-2, in the Sep. 8 experiment, the chill water valve position maintained at ~30% during the peak load, even when the room temperature continues to rise above the cooling setpoints, which implies the actual cooling capacity is only 30% of the designed capacity. This is a common practice in mitigating condensation risks. Condensation is prevented by applying the dew-point control system that maintains the chilled water

¹The two classrooms have 40 students as designed capacity

temperature at least above the dew-point temperature of the occupied space when drainage system is not available [70]. As a result, the chilled beam capacity is limited by its supply air temperature that only be cooled down to near its dew point, and maximal water mass flow rate. The reduced cooling capacity then results in a delayed system response and even cause thermal discomfort.




Figure 7-2: [Room 160] Room and chilled beam chill water valve position on Sep. 8

In this project, the slow response time and limited cooling capacity of chilled beam systems do not bring about too significant challenges because the thermal conditions of two test rooms are controlled by both the VAV and chilled beams that allow for a larger load capacity and flexibility. However, to carry out future experiments on campus, we recommend avoiding using rooms that solely rely on chilled beam systems. The nature of the system design may be incompatible with the setpoint control strategy.

7.4 Untraditional Control Sequence for Combined Systems

According to Francis Selvaggio, the Senior Building Management System Engineer from the Department of Facilities (DoF), a varying combination of HVAC room-level systems have been installed in campus buildings. The VAV with chilled beam is one of them. Other combinations include VAV with Fan Coil Unit (FCU) and fin tube radiator (FTR), VAV with FTR and so on. In some situations (Building 66 as an example), the VAV serves as a dedicated outdoor air system (DOAS) that only intakes and conditions the outdoor air and supplies it to the room level. Unlike a common DOAS that supplies a fixed amount of fresh air based on pre-designed room occupancy condition, the VAV-enabled DOAS varies the air supply volume depending on not only the occupancy, but also indoor temperature. For example, in the cooling season, during the daytime when room cooling load increases, the VAV opens up the valve and introduces more outdoor air conditioned at the central AHU that

could otherwise be replaced by recirculated air to save energy. During the nighttime, when room temperature drops below the heating setpoint, the supply air from VAV is reheated to warm up the space. To use a VAV system that supplies outdoor air to meet room setpoints is an inefficient solution because extra cooling energy is used to condition outdoor air at AHU, while extra heating energy is provided at room-level to reheat the space.

Notably, the aforementioned system combinations are indeed uncommon in mechanical design and therefore require a unique control sequence to coordinate system operations from the initial design phase. This is confirmed by Francis Selvaggio: rather than adopt standard PID control sequences, he had to manually define the specifications to ensure the systems' compatibility and controllability. Depending on the functions of the space and occupancy types, he is expected to tailor the control logics to meet the needs of space users. Those control logics may involve a combination of PID, occupancy-based, and rule-based controls that are complicated to interpret at the first sight.

Consequently, this untraditional system configuration has obtained significant attention during the commissioning phase and required careful analysis of the control sequence with the help of the illustrations on the Building Management System (BMS) dashboard, mechanical drawings, and direct conversations with the engineering and facilities management teams. More importantly, it has become a priority for the algorithm development team to consider those system operations factors when designing a load or GHG optimization model. The adjustments of the experiment methods discussed in the previous section are the results of the findings of the commissioning. To ensure more efficient execution of future pilot experiments and applicability of the proposed solutions, we recommend selecting test rooms with common system designs and control logics, and obtain a comprehensive understanding of the space before designing the algorithms and implementing experiments.

7.5 Varying Supply Air Setpoint of AHU

Per industry standard, the supply air setpoints at the AHU side are kept constant at approximately 55 °F, which has become a major assumption in the system sizing process. In a typical system or duct sizing, the lowest supply temperature helps define the maximal air flow to fulfill the peak load demand in the cooling season. The calculated maximal air flow rate, and supply air temperature, together with a safety factor (usually 150%) determines the supply capacity. MIT campus building system sizing follows the same process.

However, MIT engineers in a previous initiative looked for ways to further optimize energy efficiency for the AHU operation. After conducting rigorous psychrometric analysis on supply air setpoints, they decided to use a varying supply air setpoint schedule for certain buildings on campus, including Building 66. As is explained in Figure C-1, the supply air dry bulb setpoint is decided by the outdoor air temperature. When outdoor air is below

52 °F in dry-bulb or below 55 °F in dew-point, the supply air setpoint at AHU should be maintained at 65 °F; otherwise, 55 °F. While there's no opportunity to verify the design decision with the engineers, the hypothesis is, in doing so, the system would decrease cooling energy consumption for chilled water. Otherwise, energy on the AHU level could be wasted in the warmer or shoulder seasons when little cooling load is present, or even worse, when reheat on the room level turns on to make the supply air warmer.

While this intention is understandable, it adversely impacts the system's ability to meet room load demand. As elaborated, AHU's supply air temperature defines the system's capacity, so increasing the supply air setpoint leads to a decrease in its capacity. This results in failure to track room setpoints. As is shown in Figure 7-3, during the daytime (after 10am), the VAV supply air temperature is maintained near 65 °F because the AHU is in the energy saving mode when outdoor temperature is low. As room temperature peaks at noon, exceeding the cooling setpoint by almost 2 °F, the VAV increases supply air flowrate to its maximal rate at 700 cfm. It takes 2 hours for the room temperature to go down below the cooling setpoint. The limited cooling capacity resulted from a higher supply temperature makes the system harder to respond in a timely manner. In addition, if cooling loads were even larger with higher occupancy, the system may not be able to provide adequate cooling load and introduce thermal discomfort. All these potential issues could be prevented if shoulder seasons with a wider diurnal temperature range were considered when setting and upgrading a varying setpoint policy.



Figure 7-3: [Room 160] Setpoints, room temperature (top), VAV supply air temperature and supply air flowrates (bottom) on Nov. 9

7.6 Frequent Oscillation of Supply Air Temperature

Another issue associated with the supply air temperature is its frequent oscillation, as is shown in Figures 7-3 (bottom) and 7-4, in both heating and cooling seasons, we observed a 5-to-7 °F supply air temperature oscillation every 10 minutes independent of room equipment reheat and cooling. So far this pattern has not been found out to be correlated with any system design policy or building operation guidelines. While this problem does not directly impact research outcomes, we suggest looking into the abnormal temperature oscillation cycle.




Figure 7-4: [Room 160] VAV supply air temperature and its reheat valve position on Sep. 11. At 9:00AM - 7:00PM, supply air is not further conditioned by the chilled beam or VAV reheat coil; the measured temperature reflects the oscillation happening at the AHU side.

7.7 Sensor Placement and Calibration Susceptible to Human and Instrumental Errors

To monitor room thermal loads, a number of sensors² and meters were deployed in the two rooms and basements, including temperature sensors, flowrate meters, and occupancy sensors, summarized in Table 7.1.

Over the two-year project period, MIT's investment into purchasing, maintaining, and calibrating sensors has been significant. Unfortunately, even at present, two years into the project, we are still facing sensor installation and calibration issues that appear periodically but cannot be fully understood and fixed. For example, the two flow meters measuring hot water flows in the two rooms give drastically different readings when the rooms are in almost identical situations. The relationship between the measured flowrate of hot/cold water and

²Some experiments used illuminance sensors to record solar and room lighting conditions.

the valve position is unclear, as illustrated in Figure 7-5. A lower valve opening is not necessarily associated with a higher flowrate; zero flow does not guarantee a zero percent valve position. This uncorrelation may come from (1) sensor measurement error, or (2) the operation mechanism of valves that leads to a delayed response in controlling flow.




Figure 7-5: [Room 154] FTR valve position and hot water flowrate comparison

The sensor issues could result from both human and instrumental errors. While the former one can be prevented, the latter one is inherent to the experiment design. For example, vibration introduced by the central AHU operation could loosen meter attachment on the pipes, which could cause inaccurate readings.

Table 7.1: Sensors, meter types and their usages in onsite experiments³

Sensor	Meter Type	Usage
Temperature sensor	Platinum Resistance Thermometers (PRT)	Room temperature, VAV supply air temperature, chilled water supply and return temperatures, hot water supply and return temperatures
Flowrate sensor	U1000MKII-FM Fixed Clamp-on, Ultrasonic Flow Meter	VAV supply air volume, chilled and hot water flowrates
Occupancy sensor	Provided by Density https://www.density.io/	Room occupancy count

³Meter specifications are listed in Appendix C.

To ensure everyone's effort is well-compensated and experiment productivity improved, we recommend following three implementation steps for future experimental validation. First, select meters that have an appropriate measurement range and sensitivity level. For example, the measurement range of a flow meter should match the expected flowrate in the hot and chilled water pipes, and the sensitivity level⁴ should be adjusted to prevent both noise and information loss. Second, sensor installation and calibration should be extremely carefully treated. The installation process should follow instructions in the user manual. The positions of the sensors should be maintained in similar places, if possible, within the two rooms. The units of measurement, data post-processing methods, and data transmission channels should be examined with special attention. And finally, frequently inspecting the measured data and promptly identifying potential measurement errors will ensure that the sensors can be calibrated or replaced in time, which would minimize experiment repetition and its associated wasted time and human labor.

7.8 Inconsistent Observed Thermal Response

After adjusting the minimal VAV air flow from 450 cfm to 250 cfm in Rooms 154 and 160 (reason explained in the previous section), we cross-referenced thermal loads not only in the two test rooms, but also in Room 156, the middle room connecting 154 and 160. As is shown in the floor plan (Figure 6-1) in Chapter 6, the room size of 156 is smaller compared to the others. Given a smaller room size and its adjacency to Rooms 154 and 160, we would expect Room 156's load to be similar to, or smaller than the other two. However, we observed that only after the adjustment of minimal supply air flowrate in Rooms 154 and 160, Room 156's VAV cooling load has been consistently twice as much as those of 154 and 160. Take one baseline experiment weekend, shown in Figure 7-6, as an example. During the test weekend, setpoints are identical across the three rooms. With further investigation, we infer that Room 156's VAV load doubles because its minimal air flow is still kept at 450 cfm, nearly two times as the minimal rates of the other two rooms, while supply and return air temperatures in the three rooms remain similar.

This phenomenon raised a fundamental question about the room's thermal response. Given a similar external and internal condition, how could Room 156 have a much higher VAV load than the other two rooms? If VAV in Room 156 supplies a constant 450 cfm flow, how can its supply and return temperatures be the same as the others? With more cooling supplied to Room 156, why is its thermal response observed to be the same as the others? To figure out the answers to those questions, we need a systematic evaluation of HVAC control sequence design and examination of sensor placement and calibration, which should be a priority in future work.

⁴The minimal flowrate that can be detected by the equipment




Figure 7-6: VAV loads of Rooms 154, 156, and 160, during the baseline experiment weekend of Oct. 9 - Oct. 10

7.9 Unrobust Data Pipeline

As illustrated in Table 7.1, seven data types recorded by three types of sensors, with their associated seven independent databases are not directly managed by MIT. Specifically, BTU Meter, VAV Energy, Occupancy, and Setpoint Control are piped through platforms developed by Schneider Electric, contracted by MIT to install and maintain BMS in many campus buildings. Grid Carbon Intensity and Weather data points rely on external parties⁵. Only the Thermal Comfort Survey is fully developed and maintained by our research group.

Seeking professional services via industry partnership allows for a systematic approach to manage campus buildings, perform fault detections, and regularly maintain the buildings. Indeed, it is one of the very few options universities have for campus building management in recognition of the challenges of working with large-scale building operations. However, the cross-team communication and collaboration inevitably prevented us from updating the system that is integrated with functions developed for the experiments in a timely manner. It took us more than a year to fully establish this data pipeline, a period that included frequent database breakdowns, follow-ups with the engineering team to diagnose issues, and numerous debugging procedures. As a result, the team missed ideal testing periods in the cooling season due to data logging and transmission issues as well as inefficiencies in the project management, and only started experiments after 1.5 years into the project.

To prevent a similar lengthy setup process from happening again, it is recommended to plan ahead, get an overview of the ecosystem of the MIT BMS, and efficiently communicate with all the stakeholders to ensure the project goal and motivation are aligned. MIT is also encouraged to develop an automated database that could collect and process different data streams, which will drastically save manual efforts throughout the entire project execution.

⁵MIT's own weather station does not forecast or collect solar data that is necessary in the experiment.

Chapter 8

Experiment Results and Discussion

Chapter 6 points out the two phases of onsite experiments: the first phase establishes a baseline load for the optimal control test, and the second phase performs the optimal control test and estimates the savings relative to the baseline. It also suggests a simple rule-based control algorithm as a benchmark for the machine-learning based control as an additional method comparison. This chapter first addresses the setup of the two phases and configuration of a rule-based control algorithm, followed by the critique of baseline validation methods, and concludes with the estimated savings of energy and GHG emission from optimal control days with respect to the baseline and rule-based control. For the rest of the chapter, baseline refers to the days implemented with a constant 24-hour heating and cooling setpoint schedule; optimal control refers to the active test days with optimal setpoints output by the two-step optimization framework; and RB benchmark refers to the days with a simple rule-based setpoint control.

8.1 Baseline, Optimal Control, and Rule-based (RB) Benchmark Setup

Figure 8-1 below illustrates actual room temperature, cooling setpoints, and heating setpoints from a baseline day (Nov. 25) and optimal control day (Nov. 26) in a heating season. All data points are the real time measurement collected from the thermostats in the rooms and systems and recorded in the Clockworks database. The first baseline day deploys a constant heating setpoint at 69 °F and cooling setpoint at 72 °F¹ The second optimal control day uses a constant cooling setpoint at 76 °F and a varying heating setpoint schedule. The 76 °F cooling setpoint is pre-determined, and can be tuned based on results of the thermal comfort survey. The heating setpoint schedule is the output of the optimization algorithm.

¹The baseline setpoint schedule may be slightly different in each experiment due to hardware and software configuration, but typically the variation is within 0.5 °F.

In the baseline, the actual room temperature oscillates around 69 °F except for the 8-10am period when solar heat gain introduces cooling loads. In the optimal control, room temperature does not exactly follow the heating setpoint as heating setpoints ramp down rapidly during the nighttime unoccupied period. This is because the room temperature then naturally decays, driven by the conductive heat loss with ambient condition and the convective heat loss through infiltration and minimal fresh air requirement (250cfm) from the VAV system². When room temperature hits an upward heating setpoint trend in the early morning, the heating system modulates on, and the actual room temperature matches the desired heating setpoint. This validates the heating load increment from the designed heating setpoint ramp-up rate³ is achievable with the capacity of room heating systems (FTR and VAV reheat). Further tests on the system capacity will allow a larger setpoint step change, which gives more flexibility and opportunities for setpoint schedule design and GHG savings.




Figure 8-1: [Room 154] Room temperature, heating setpoints, and cooling setpoints for baseline (Nov. 25) and optimal control (Nov. 26) days

The RB benchmark test involves a varying heating and cooling setpoint schedule, only based on the occupancy information⁴. During the unoccupied period, heating and cooling setpoints are set at 64 °F and 76 °F respectively; during the occupied time, the range is narrowed down to 69 °F and 72 °F, the same as the baseline test setpoints. Unlike the optimal control algorithm that involves a constraint on setpoint change rate, the benchmark model only adjusts setpoints immediately after the schedule changes, which leaves little time for system response. As is observed in Figure 8-2, setpoint adjustments occur at 9am, after which the system takes 30 minutes to reach the target room temperature (from 66.8 °F to

²The VAV provides a constant 250 cfm throughout the day to fulfill indoor air quality standard. The minimum flowrate in all experiments are not determined by the actual room schedule.

³Limited to 0.5 °C/ 0.8 °F over an hour, as a constraint in optimization.

⁴All the occupancy schedule used in baseline, optimal control, and RB benchmark experiments are identical. Rooms are occupied 9am-2pm, and 3pm-6pm.

69.0 °F).




Figure 8-2: [Room 154] Room temperature, heating setpoints, and cooling setpoints for RB benchmark (Nov. 29)

8.1.1 Load Calculation Example

Since the majority of experiments were performed in the heating season and heating load calculation involves more intricacies than cooling load calculation in the cooling season, this section provides an example of load calculation in the heating season first.

In heating seasons, load calculation for FTR and chilled beams is straightforward⁵. However, VAV loads need to exclude the cooling or heating load that outdoor air brings in for air quality purposes rather than to maintain setpoints, which is unique to the system configuration of the two test rooms. To illustrate, Figure 8-3 identifies the VAV airflow actively contributing to room heating (in light blue) and cooling loads (in green), with remaining flows labeled in yellow. Depending on the supply air temperature and room temperature, the indoor flows can bring either heating or cooling loads into the room disregarding if rooms actually have cooling or heating demand. For this specific experiment, the outdoor air is first heated up by AHU to approximately 65 °F before being supplied into the room. If the room temperature is higher than 65 °F, the air takes away heat from the room, and vice versa. These heating and cooling loads are not considered as active loads for maintaining room temperature, and therefore should be excluded in the thermal load calculations. To exclude those loads, a filtering method (elaborated in Appendix B) is applied to capture all effective airflows based on the variations in the actual supply flowrates, setpoint changes, and observed room temperature movements. Albeit with minor errors due to the complex relationship among the observed data points, this filtering logic allows for a relatively precise estimation of the cooling and heating loads used in the experiment result analysis.

⁵The presence of cooling loads in heating seasons is due to the dual setpoint temperature control.




Figure 8-3: [Room 154] An example of effective VAV airflow for heating and cooling

Figure 8-4 illustrates the calculated heating and cooling loads from the FTR, chilled beam, and VAV, along with the ambient temperature. We can make a few conclusions about the thermal response in Room 154. First, the majority of heating loads are provided by the FTR, with the maximum load on the order of 10,000 BTU/h. Second, the timing of the peaks of heating load from the FTR match those from the VAV reheat, which is consistent with the sequential control logic. Third, cooling loads could occur even during the heating season especially when massive solar heat gain is present. For example, the large solar load during the daytime (8am-12pm) of Nov. 25 leads to room temperature increase, drives the heating loads to zero, and triggers both chilled beam and VAV cooling.




Figure 8-4: [Room 154] Heating and cooling load profiles from FTR, chilled beam, and VAV, and outdoor temperature for baseline (Nov.25). Heating loads are positive and cooling loads are negative.

In the cooling seasons, load calculations for all types of equipment are easier because the conditioned outdoor air supplied by AHU not only serves to meet air quality requirements but also provides space cooling.


8.2 Comparison of Baseline Estimation Methods

Four methods are proposed in Chapter 6 for establishing a reasonable baseline for Room 154 when optimal control test is conducted. They are: (1) a parallel test, (2) a sequential test, (3) a scaled predicted baseline, and (4) a machine-learning based baseline. With more than eight sets of experiments in heating, cooling, and shoulder seasons, we recommend the third approach – using a scaled predicted baseline from Room 160 – because it imposes fewer implementation hurdles and can provide the most reliable baseline. The rest of the section evaluates the strengths and weaknesses of each method with experiments performed across cooling, shoulder, and heating seasons.


8.2.1 Validation 1: Parallel Test

Due to differences in room solar heat gains, uncontrollable occupant behaviors, and sensor measurement accuracy, we conclude that it is extremely challenging to use Room 160 as a perfect baseline for Room 154. The following comparisons are conducted on three types of loads: VAV heating and cooling, FTR heating, and chilled beam.

The VAV loads in the two rooms demonstrate above 70% load profile proximity using Pearson correlation between two load datasets. The correlation is higher during the night-time when solar radiation and occupancy effects become less prominent. Figures 8-5a and 8-5b illustrate load profile comparisons only considering the VAV system in the cooling and shoulder seasons. In Figure 8-5a, both loads peak at a similar magnitude. The base load of Room 160 during the non-peak hours is consistently ~50% lower than that of Room 154 while following the same trend. As Aug.18 and 19 is the first test in the series, except for room access restriction, other external factors were not rigorously controlled, including blind, lighting, projector, and connecting door status. To understand if the VAV load differences are due to sensor errors or the external environmental factors, a more rigorous control over room conditions is deemed necessary.



(a)



(b)

Figure 8-5: Baseline cooling load (only VAV) comparison in Rooms 154 and 160 on Aug. 18 and Aug. 19 (a) in the cooling season, and on Oct. 16 and Oct. 17 (b) in a shoulder season. Cooling load is positive, and heating load is negative⁶.

In the baseline experiment on Oct.16 and Oct. 17, we put signs in the rooms, requesting students not to use the space, or to follow the instructions when they study to ensure good experiment conditions. We also visited the site to check room conditions during the two weekend days. Those control strategies overall managed to improve the room load similarity, as observed in Figure 8-5b. Despite the fluctuations in Room 160’s load during the midnight of Oct. 16, the rest of loads are either almost identical or different for explainable reasons. In our further analysis of multiple baseline comparisons, we identify that Room 160 consistently has a higher peak load than Room 154 does, especially in the shoulder season. While most of variables are controlled in the follow-up tests, one important difference between the two rooms is their location within the building and external environment, which affects the

⁶Detailed experiment documentation can be found in Appendix B.

room's solar heat gain.

To verify this hypothesis, we first walked around the building and visually inspected the surroundings. From the below four photos (Figure 8-6) taken from the interior and exterior of the rooms, it is obvious that in front of Room 154 stands a taller tree than Room 160, creating more shade during the middle of the day. It partially explains the difference in noon loads observed in October. Nonetheless, the solar condition during the cooling season may be different as the baseline analysis on Aug. 18-19 shows Room 160 had a lower load compared to Room 154. Hence we can further hypothesize that the difference between solar irradiance varies across seasons because of the varying solar path within a year. A model can be developed to simulate the annual solar condition of the rooms and generalize the patterns.

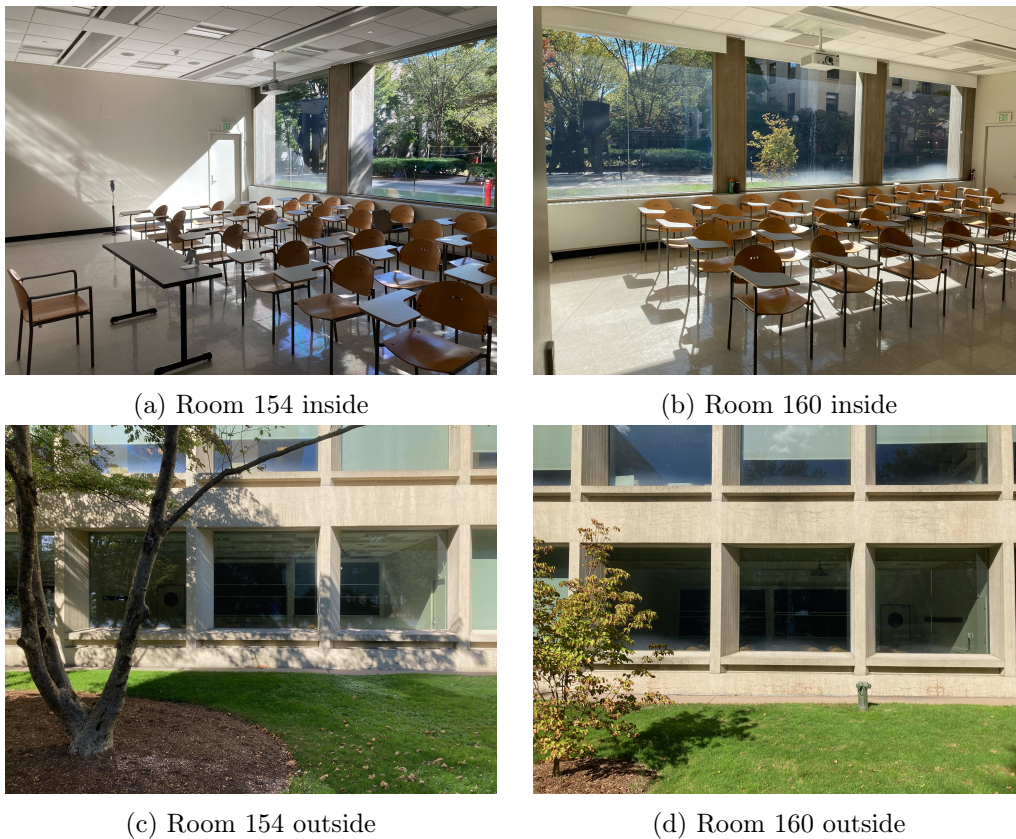


Figure 8-6: Inside and outside conditions of Rooms 154 and 160⁷

Aside from visual inspection, illuminance and temperature sensors were deployed in both rooms during the weekends on Nov. 20 - Nov. 21 and Nov. 25 - Nov. 27.⁸ The light intensity data conform to the observed ambient environment and the shading effect. Among the three-day test on Nov. 25 - Nov. 27 (shown in Figure 8-7), the second day is a cloudy

⁷All photos were taken on Oct. 19 noon, credit to Dr. Kevin Kircher.

⁸Sensor positions are described in Appendix A.

day, so sunlight coming into both rooms is scarce. On the other two sunny days, Room 154 has a consistently lower average and peak light intensity. Additionally, the light intensity of Room 154 dips during the morning period (from 9-11am) across all experiment days, while the same behavior is not observed in the other room. Based on light intensity data, it can be concluded that the actual solar irradiance incident on the glass panels is different due to external shading and room locations, creating different solar loads in the two rooms especially on sunny days.






Figure 8-7: Light intensity comparison of Room 154 and Room 160 measured from the middle glass panel on Nov. 25 - Nov. 27

The difference in solar heat gains not only impacts VAV loads but also chilled beam loads in summer and FTR loads in winter. However, solar heat gain is not the determining factor in explaining differences in chilled beam and FTR loads across the rooms. As is described in Chapter 7, the flowmeter installation and calibration have led to drastically distinctive measurements of hot water flows passing through the FTR, making FTR loads in Room 160 almost as five times smaller as loads in Room 154, as is presented in the flowrate measurements for the two rooms on Nov. 25 - Nov. 27 in Figures 8-8a and 8-8b. During the three-day period, Room 154 has a peak FTR hot water flow above 1 gpm (gallon per minute), while Room 160's only reaches 0.04 gpm, a factor of 25 lower. While the supply and return water temperature difference is larger in Room 160, the total heating load difference is still dominated by the flowrate gap. Investigating the valve positions of both FTRs demonstrates very similar average opening positions in the two rooms. Two hypotheses could be made based on the observations. First, the flowmeter readings in one of the rooms (Room 160 more likely) suffers from inaccuracies due to sensor position and its inherent sensing mechanism. From the Schneider Engineering team we learned that "if the rooms aren't piped identical then the readings will differ between the two preferring the one with lower back pressure." Second, the readings for two rooms are both correct, but

the equipment control and commissioning are different. Because the two rooms have two heating systems (VAV reheat and FTR), evaluating their heating loads solely from the FTR is biased. Especially because we understand the two systems complement each other – when load provided by one system cannot meet the demand, the other system would respond to fulfill the load gap – we shall look at the heating loads from VAV reheat to gain a holistic picture. We would expect that, if the two rooms indeed have comparable loads and loads from the FTR in Room 160 is significantly lower than that in Room 154, then the heating loads of VAV reheat in Room 160 should exceed those in Room 154. The result analysis section shows this is observed in VAV reheat, with the exception that the VAV reheat load in Room 160 is only twice as high as at in Room 154. Hence the total heating loads of Room 154 are still higher than those of Room 160, but in a more reasonable range.



(a)



(b)

Figure 8-8: Measured hot water flowrates through FTR in Rooms 154 (a) and 160 (b) on Nov. 25 - Nov. 27 (heating season)

In addition to solar heat gain, the occupant behaviors also could contribute to significant

room load behavior difference. Analyzing weekday data when different class schedules take place, we find that room loads are largely driven by how occupants used the space. For example, a large occupancy leads to a large heat injection; students can raise or lower the blinds to create a visually comfortable environment; and lecturers can leave the projectors running after leaving the classrooms. All of those factors can disrupt a baseline model in providing the benchmark for the other room.

In brief, using a twin room as baseline reference needs further consideration and investigation. Although the two rooms share very identical room and system configurations, the subtle difference in room location and room usage pattern has given rise to difficulty in establishing a reliable and accurate parallel baseline. While this method may be able to provide a straightforward evaluation, it is susceptible to large-scale implementation difficulties because of the unlikelihood of identifying a pair of rooms to establish a reference for every experiment site.

8.2.2 Validation 2: Sequential Test

A sequential baseline method uses Room 154's first experiment day as its own baseline for the following optimal control day to escape the challenges of establishing two identical test rooms suggested in the parallel test. However, using the first day experiment as a true baseline for the second day, even within the same room, is also demanding, due to the salient load disturbance factors such as outdoor temperature and solar intensity.

Take the experiments during Nov. 25 and 26 as an example (in Figure 8-9), where the first day is intended to serve as a true baseline for the second optimal control day. The measured solar intensity in Figure 8-7 and real-time outdoor temperature in Figure 8-10 all indicate that the increase in room temperature from 8am to 10am of Nov. 25 is mainly due to the solar heat gain. This leads to a spike in cooling load and warms up the space for the rest of the day. In contrast, the following cloudy day has very little direct solar irradiation, so the heating load is evenly distributed throughout the day.

Differences in solar exposure during the two test days make it challenging to accurately identify and estimate load savings. For instance, while the second day does not involve cooling loads, it is not a result of a raised cooling setpoint (as part of our optimal control strategy), but because there is indeed little solar load. Likewise, we cannot claim the optimal setpoint control causes a higher peak load during 12-3pm (nearly 17,000 BTU/h) than the baseline test with a peak load around 10,000 BTU/h in the same period. This is because on the baseline day the construction materials of the room store thermal energy from the sun and start to release heat as the surrounding environment becomes cooler, while on the optimal control day the thermal storage capacity is not fully utilized, and the room needs additional energy to heat up the space as the ambient temperature decreases.

Not only by solar load, establishing a baseline load can also be influenced by outdoor

temperature. According to thermal transfer principals, the rate of heat transfer is related to the temperature difference between two materials. The larger the difference, the higher the transfer rate. In this case, considering the average indoor temperature almost constant, the average outdoor temperature during the two test days affects the average heat exchange rate. A colder weather leads to a faster heat loss and therefore a higher demand of heating loads. Unfortunately, the changeable local weather has never allowed us to find the identical weather conditions in four sets of experiments, which signifies the difficulty of obtaining an undisturbed baseline.






Figure 8-9: [Room 154] (Same as Figure 8-1) Room temperature, heating setpoints, and cooling setpoints for baseline (Nov. 25) and optimal control (Nov. 26) days

Consequently, our analysis capacity is limited to some small time intervals (such as on the order of a couple of hours) when most factors are under control and aligned. For instance, at 12am – 6am of Nov. 25 when room temperature is maintained at 69 °F and outdoor temperature at around 35 °F, the heating load peaks at 24,000 BTU/h; on the other hand, at 6pm – 12am of Nov. 26 when room temperature drops below 69 °F while the outdoor temperature is also close to 35 °F, the peak load decreases to 12,500 BTU/h, with a similar heating frequency as the previous test day. Such a 6-hour horizon analysis showcases the success of optimal setpoint control in reducing heating loads. However, the sequential test method is rather inefficient and unproductive, and can be subject to unpredictable experiment constraints.



(a)



(b)

Figure 8-10: Heating and cooling load profiles of FTR, chilled beam, and VAV, and outdoor temperature for Room 154 baseline on Nov. 25 (a) (Same as Figure 8-4) and optimal control on Nov. 26 (b)


8.2.3 Validation 3: Scaled Predicted Baseline

To decouple and minimize the effects brought by room properties and external environments, we integrate the advantages of the first two approaches and design a predictive baseline of Room 154 based on the consistency in load behaviors between two rooms. In the two experiment sets on Dec. 26 - Dec. 30 and Jan. 1 - Jan. 3, we conducted two baseline, control, and RB benchmark tests in Room 154, and four baseline and two RB benchmark tests in Room 160. Although we realize from the parallel test that the two room loads differ from each other, if the difference is consistent, for example, by a constant factor or can be easily correlated and interpolated, then we can predict Room 154's baseline by adjusting Room 160's load on the day of Room 154's optimal control.


All heating and cooling loads for the baseline and RB benchmark experiments are calculated for a 23-hour period to exclude the load transitioning period. Because the cooling load size is negligible in the testing period (heating season), we focus on the heating load from the VAV reheat and FTR. As illustrated in Figures 8-11a and 8-11b, the linear correlation of the 23-hour total heating loads across the two rooms are significant with R^2 values exceeding 0.999. Although this correlation test is conducted on four test days that may not generalize future load data well, the high R^2 values ensure the applicability of the model to the optimal control days within the test period. In the meantime, we also recognize the actual load relationship may not be linear, because of the differences in thermal response to solar load, room occupancy status, and other external environmental factors. To illustrate, our investigation suggests in winter Room 160 receives more solar load than Room 154 does. A change in solar heat gain may lead to a non-linear heating and cooling load changes, and therefore introducing non-linearity in 23-hour load sum.

Although we launched eight experiments from 2021 summer to 2021 winter, only the last two sets of experiment data can be applied to this analysis. The constant tuning and calibration of sensor measurements lead to different load data distribution from each individual experiment, which makes them incomparable with each other. To further validate this hypothesis and establish a more precise connection between the two load behaviors, more experiments with a consistent measurement standard are required in all seasons and weather conditions.

In the following section, we adopt this method to estimate the energy and GHG savings by adjusting the baselines of Room 160 using the coefficients defined by the regression equations.



(a)



(b)

Figure 8-11: Correlation of HVAC reheat load (a) and FTR load (b) between Room 160 and Room 154

8.2.4 Validation 4: Machine Learning Driven Baseline

The last approach is an extension of the third one. Instead of using Room 160's load as the only feature for a predictive model, this method could potentially integrate all necessary information available to form a machine learning driven baseline for Room 154. With these features and advanced learning algorithms, the model is expected to capture all the subtle load changes caused by solar, temperature, occupancy, or system-level conditions. However, training an accurate learning model requires a high-quality and rich dataset that takes time to gather. As discussed later in Chapter 9, developing a small-scale predictive model is more difficult and needs a larger dataset than training a large-scale model. A comprehensive data collection protocol shall be developed with future research effort to specify the type of data, collection period, and load perturbation schedule (to provide rich information on room's thermal dynamics).

8.3 Optimized Setpoint Control Results




Figure 8-12: Comparisons of the average electricity carbon intensity for baseline and optimal control tests

While the ultimate optimization target is to minimize GHG emissions, the following analysis is conducted based on load savings, as we find the changes in total carbon emissions are mostly driven by changes in thermal loads. Figure 8-12 compares the weighted electricity carbon intensity of baseline and optimal control tests conducted on Nov. 20 - 21 and Nov. 25 - Nov. 27. The weighted average of carbon intensity is calculated by dividing the total carbon emission by total electricity consumed in a designated time period. The weighted averages on Nov. 20 and Nov. 21 tests are performed on a 7-hour period. Two optimal control experiments are conducted on Nov. 26 and Nov. 27 following a baseline test on Nov. 25. Weighted averages on Nov. 25 - Nov. 27 tests are calculated on a 12-hour

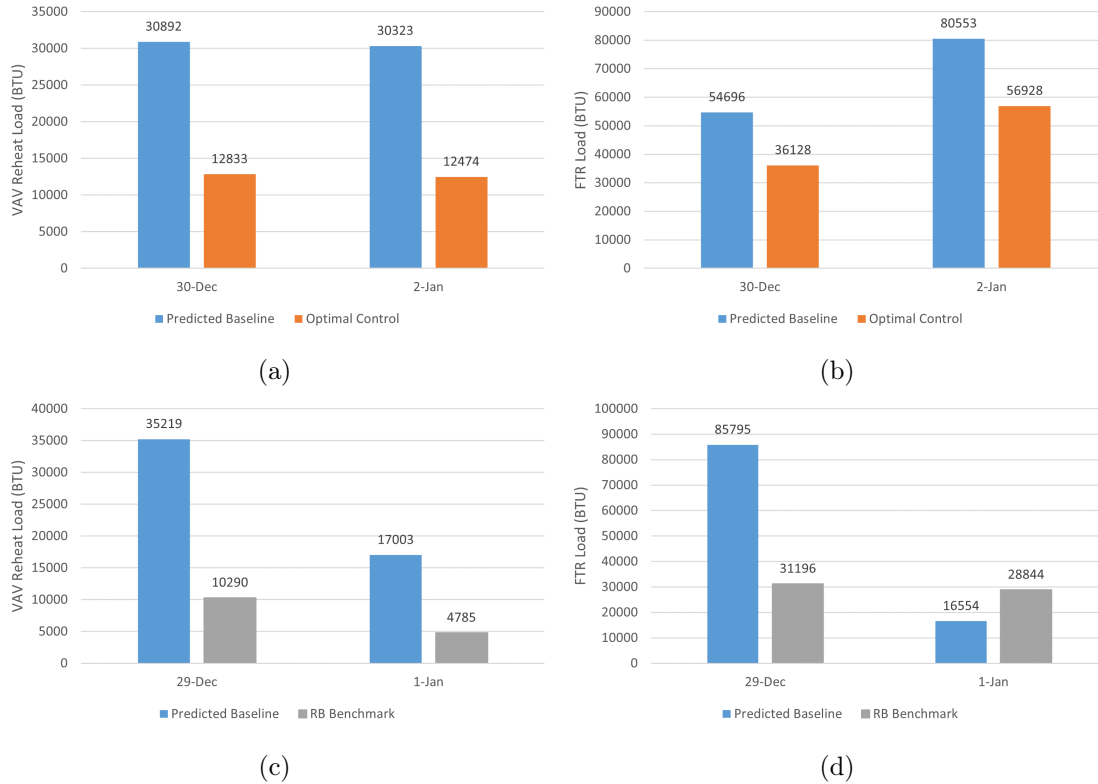


Figure 8-13: [Room 154] VAV reheat and FTR load comparison among predicted baseline, optimal control, and RB benchmark⁹

interval without solar exposure and 23-hour interval with solar exposure. The comparisons demonstrate no significant load shifting behaviors as the weighted average carbon emissions in the optimal control tests are very close to those in baseline tests. Two reasons can explain this phenomenon. First, the daily grid carbon intensity profile fed into the optimization algorithm is the same across all experiment periods, which creates little input variations in the optimization problem. Second, this phenomenon is aligned with simulation results in Chapter 4, which verifies the underlying optimization mechanism in this framework.

8.3.1 Load Comparison

Taking the slopes of linear regressions defined in Figures 8-11a and 8-11b, we scale Room 160's load on the baseline days to estimate the baseline loads of Room 154¹⁰. We then can estimate the load savings from the optimal control days and RB benchmark days. Figure 8-13 compares the loads of three types of experiments on Dec. 26 - Dec. 30¹¹, and Jan. 1 - Jan. 3. Heating loads are separately analyzed for each system type because the meter installation and calibration for the FTR and VAV may be different. Combining two heat

⁹ All comparisons are based on 23-hour periods.

¹⁰ Step-by-step calculation is included in Appendix B.

¹¹ No experiment data were collected on Dec. 27 due to the experiment errors.

sources may lead to loss of critical information about the individual system. Note that all the Predicted Baselines are not the actual observed baseline load values because those are the days when Optimal Control and RB Benchmark experiments are performed.

On average, for Room 154, the optimal control days manage to save 30% FTR loads and 60% VAV reheat loads; the RB-benchmark days achieve a higher saving of 70% FTR loads and 70% VAV reheat loads. (The RB Benchmark for VAV reheat on Jan. 1 (Figure 8-13c) is higher than the Predicted Baseline. We believe the unreasonably low Predicted Baseline is a result of prediction inaccuracy¹², not a reflection of the reality, and therefore should be treated as an outlier.) This is expected because the RB-benchmark would only adjust setpoints according to the pre-determined occupancy schedule without pre-conditioning the room, which contributes to the additional energy savings.




Figure 8-14: Distribution of temperature deviations of two baseline, optimal control, and RB-benchmark experiments

However, the extra energy savings for RB benchmark come with the sacrifice of thermal comfort. Figure 8-14 illustrates the distribution of room temperature deviations from the desired range in the occupied hours¹³ for three types of experiments. The data points in the box-plots are temperature deviations measured and recorded on a 5-minute interval. The sum of temperature deviations based on a 46-hour period (in Table 8.1) is the highest in the baseline test, followed by RB benchmark, and lastly the optimal control test. So are the average temperature deviations. In addition, the distribution of deviations also differs. The baseline tests incur fewer extreme deviations than the other two because the actual room temperature is maintained narrowly between a 3 °F deadband without large fluctuations. Optimal control tests, while introducing a wider window of room setpoints and temperature, prevent temperature fluctuations and minimize thermal discomfort by pre-conditioning the

¹²The outdoor temperature on Jan. 1 was higher than any other testing days. The baseline of 160 is predicted by the outdoor temperature with a UA value (discussed in the later section), and is further used to predict the baseline of Room 154. So the cumulative prediction errors especially during warmer ambient condition lead to an unreasonably extremely low baseline heating loads.

¹³The figure only considers non-zero values for the two test days.

Table 8.1: Evaluation of thermal comfort for baseline, optimal control, and RB-benchmark test days

	Baseline	Optimal	RB Benchmark
Sum of temperature deviations (°F)	-46.91	-35.87	-40.23
Average temperature deviations (°F)	-0.49	-0.41	-0.46

space. On the other hand, the Rb-benchmark tests, while achieving the greatest savings, cause larger temperature swings during the beginning of the day (approximately 30 minutes before the first class of the day). By comparing heating loads and thermal comfort level, we detect a trade-off between these two targets; a data-driven robust predicative optimization framework provides us with a toolkit to achieve a balance between the two objectives.

8.3.2 Relationship Between Room Loads and Outdoor and Indoor Temperatures

Based on thermal load calculation, the room load is proportional to the difference between inside and outside temperatures, with conductive heat flow through the outside wall and window and convective heat flow associated with heating outdoor air brought into the building. Adjusting heating setpoints based on the occupancy schedule manages to lower the average room temperature and therefore reduce the temperature difference in a heating season. Figure 8-15 compares room average temperature of baseline and optimal control days in four experiments, including both Room 154 and Room 160. We observe that average room temperatures during optimal control tests are consistently 1 to 1.5 °F lower than those from the baseline tests, which proves the feasibility of lowering room temperature with setpoint control.

In the meantime, average room temperature is determined by not only the setpoints, but also the ambient environment. To be specific, most of the baseline, optimal control, and RB-benchmark experiments have almost identical setpoint schedules (Figures 8-1 and 8-2 showcases the setpoint schedules for baseline, optimal control, and benchmark as an example), but the actual room temperature varies because a colder outdoor condition increases the rate of heat loss with interior. To account for this factor, we plot the trend between outdoor temperature and heating loads of the four baseline tests in Room 160 in Figure 8-16. Identical setpoint schedules are implemented in the four baseline tests, so the heating load trends can more accurately capture the impact of outdoor temperature on the heating loads. A linear trend can be observed between the two variables, and the slope can be used to experimentally determine the UA value of the space. This provides critical information about the room’s heat transfer in response to outdoor temperature changes.

Based on our limited knowledge of room’s construction and insulation materials, the estimated UA value of Room 160 using R-10 wall and single-pane glass properties is 340




Figure 8-15: [Room 154] Average room temperature in baseline tests and optimal control tests




Figure 8-16: [Room 160] Outdoor temperature vs. heating loads

BTU/h-°F. The ballpark estimate is very close to the observed value (235.8 BTU/h-°F) through experiment, which gives some level of confidence that the room is properly controlled and results are not skewed. However, it is important to note that the UA value from the experiments includes not only conductive heat transfer but also convective heating loads introduced by the VAV outdoor air. While a typical construction has a constant UA value determined by the material properties, the test room may not follow this trend. We envision

a constant UA value within a certain range of outdoor temperature, beyond which room heat loss is dominated by the VAV outdoor air. For example, as the outdoor temperature increases in the heating season, the rate of heating load decrease will flatten and the heating load will approximate to a non-zero value defined by the heat loss incurred by VAV outdoor air. This could explain the abnormally low predicted load on Jan. 1 in Figure 8-13c when the average outdoor temperature was 48 °F, outside of the temperature range of the regression analysis.

Therefore, we recognize the need for doing more controlled tests in both rooms to improve the credibility and reliability of this important property and validate the above hypothesis, and confirming room thermal properties by accessing detailed construction information. A reliable UA value can facilitate future experiment result analysis by (1) providing support for the basis of load calculation and (2) generalizing the room’s heat transfer properties for more scaled-up experiments.

8.4 Future Work

Having witnessed all the challenges in setting up and managing experiments, we manage to identify energy and GHG emission savings from onsite experiments with the 2-step optimization framework. This over half-a-year experiment provides valuable insights into onsite experiment preparation, hardware and software troubleshooting, effective communications with various stakeholders, and research method design iteration based on real onsite conditions. Considering practicality and feasibility, we recommend future researchers to adopt a scaled predictive baseline for Room 154 based on Room 160’s load. The research team is encouraged to follow or update, if necessary, experiment protocols for conducting baseline-optimal RB-benchmark series of experiments, to ensure the validity and replicability of all experiment results. In addition, while the experiment results are promising, the team should pay attention to the remaining system- and control-level issues identified in Chapter 7.

Nonetheless, the linear-scaled baseline method is limited to small-scale implementation because finding similar experimental rooms for larger space is never easy. In that regard, we consider a machine learning driven baseline model would better fit the test purpose and can be integrated to the next step of the project. A predictive baseline model that could simulate what would have happened if the optimal setpoints were not executed better defines the research results. On top of the standard features for an energy forecast model, such as historical energy consumption rate, ambient weather condition, and temporal information, the predictive baseline model for Room 154 can also utilize load data from Room 160 that contains important information regarding the campus operations and student activities. So the next step also includes developing a predictive model and testing out its capability of “backward forecast”, and integrating it into the rest of the experiment framework.

Chapter 9

Methods Comparison and Scalability and Implementation Discussion

With constraints on project timeline and resources from MIT personnel for onsite test execution, the previous chapters present three optimization frameworks and simulated and experimented in mostly single zone conditions as a proof of concept. The three methods in summary are: (1) a three-step predict-perturb-optimize framework, (2) a two-step predict-optimize framework (a variant of the first approach), and (3) a model predictive control method. As the ultimate goal is to scale up to campus- and grid-level application, this chapter compares and contrasts the three approaches in their performance, modeling complexity, data requirement, and scalability. In addition, one critical element in two machine learning enabled methods is the high demand of prediction accuracy that is sensitive to the scale of data (for example, room level vs. building level vs. campus level). This chapter also addresses sensitivity analysis of machine learning model performance to gauge its impact on implementation on a different scale. Finally, one uniqueness of MIT campus GHG emission initiative is it requires a comprehensive evaluation of the energy supply and demand system at MIT. To optimize for the best GHG emission and energy consumption scenario, not only do researchers consider the thermal dynamics and occupant feedback, but also the compatibility of algorithms with existing infrastructures on a system level. The perspective shared in this chapter serves as a grain of thought to invite future research teams to dive into further.

9.1 Methods Comparison

9.1.1 Performance

While all three approaches demonstrate the capability to minimize GHG emission or peak load, the MPC approach fully demonstrates load shifting strategies by pre-conditioning

space, followed by the three-step method, while the two-step method does not apply load shifting at all. This can be explained from two perspectives. First, the MPC is modeled and simulated on a five-minute interval, more granularly than the other two models. The granularity of modeling leads to its increased sensitivity to changes of external factors. On the other hand, the other two models simulated on an hourly basis are less reactive due to the smoothing effect. Second, the optimization model for MPC is embedded with a 1R1C numerical model, which offers a relatively accurate prediction of room thermal performance compared to surrogate models used in the other two optimizations. The optimized setpoint control is heavily dependent on the amount of information and the confidence of information the algorithm is aware of for the predicted horizon. Because machine learning models will never produce 100% correct predictions, the optimality of control decisions can be affected by prediction errors. Nonetheless, in reality, obtaining a resistor and a capacitor value for establishing an MPC model requires some estimation (i.e. a surrogate model) and can never be perfectly matched with the ground truth.

Between the three- and two- step approaches, the three-step model is more responsive and sensitive to changes of modeling environments because it decouples perturbation from baseline phase, which allows for an efficient system identification process. It also ensures that a baseline surrogate model can pick up load trends easily from external signals following a more predictable unperturbed setpoint schedule. Admittedly, the three-step approach may suffer from loss of accuracy because errors from two separate learning models could accumulate. The two-step approach, originally designed to minimize the error accumulation, unfortunately has shown limited performance improvement. This is mainly because the enormous amount of non-linear information and trends in the two-step approach makes it strenuous for a linear model to understand. Because the algorithm is insensitive to external environment variable changes, in both simulation and the onsite experiment sessions, we observe the setpoints it outputs look very similar when outdoor temperatures are different. In addition, since the optimization is conducted on a perturbed setpoint and load dataset, its control decisions are impacted by the simulated perturbed dataset which is not aligned with the actual load trend.

9.1.2 Modeling Complexity and Scalability

Modeling complexity largely stems from two factors: (1) creating room simulation models that inform the optimization, and (2) establishing machine learning surrogate models. For a real-world implementation, the three approaches involve both factors on different levels. The two- and three- step approaches involve one or more surrogate models to generalize the simulated loads with their associated features. The MPC approach needs to first learn the thermal performance of a space and obtain reasonable R and C values for the 1R1C network model by fitting a model to the observed loads.

The two simulation methods adopted in the three approaches are (1) numerical models (in MPC and three-step method), and (2) software simulation in EnergyPlus. Numerical models can be relatively easily established with mathematical formulation and simplified modeling assumptions; in this thesis, the two numerical models are generated in Python and MATLAB. The aggregate-level model in Python enables a convenient scale-up for simulation of hundreds and thousands of zones in just a few minutes with the current coding framework. However, the EnergyPlus involves more computational complexity. The detailed modeling approach inevitably requires initialization of hundreds of parameters for a single zone simulation, based on specifications and information from building construction, design, and engineering. The parameter calibration process could even take more effort and time. If expanded to campus and grid scale, the number of parameters and computation complexity of modeling would explode, which renders the method infeasible. Therefore, we conclude that EnergyPlus or other simulations enabled by state-of-art software is feasible for small testbed for proof-of-concept, while numerical models, albeit with less accuracy, provide efficient solutions for large-scale implementation. Further exploration of other more advanced numerical models (such as the 2R2C network model) could allow us to balance the trade-offs between model accuracy and computational complexity.

9.1.3 Data Requirement

Different methods and components of models have different data requirements, which can be easy or difficult to fulfill in real life. The baseline learning model of the three-step approach only requires gathering aggregate load data from a multi-zone building or a grid. This can be easily accomplished by the existing power meters installed in the building or a large aggregate. For MPC, to identify the appropriate R and C values, room temperature data during a specified period need to be collected. For example, nights seem appropriate for a decay test to get a value for C. Occupancy data, if available, can help check for load disturbances. This can be challenging from an implementation standpoint as the researchers would need to access the occupancy data and their activities to precisely filter out the data of interest.

If we consider gathering baseline data as observation and process of data readily available in the current system, then collecting perturbation data involved in the three-step approach needs active interference of the space. Obtaining the consensus from all occupants would not be easy, and even with building owners' approval, its consequences and implications need to be well thought-out.

In addition to the accessibility of datasets, their size also influences the prediction quality. To train an equivalently good model for a single zone or a small community requires more high-quality data than that for a large aggregate. The trend of small-scale loads can be easily disturbed by stochastic occupant behaviors, whereas those outliers are averaged out

in a larger community. As a result, a future scaled-up implementation is in favor of an easier-to-train and more accurate prediction model.

9.2 Machine Learning Performance Sensitivity Analysis on Data Scale

Closely related to the data size requirement is the model’s sensitivity to data scale. Large-scale energy consumption (i.e. energy consumption on the order of magnitude of 100 MWh) tends to be easier to forecast because the noise internal to the dataset is canceled out and the trend is smoothed. Previous research efforts have focused on model performance of large-scale forecasts, but not many have quantitatively explored the impact of data scale itself on model performance. This often-neglected question is indeed significant to the MIT campus, where loads are much larger than a small condo but much smaller than the grid. Therefore, in collaboration with my colleagues, Qianqian Wan, Anne Qingyang Liu, and Will Atkinson, we conduct a sensitivity analysis on model performance at different data scales, from a single building to a larger group of 5-15 buildings.

9.2.1 Data

We work with cooling load time series data from MIT Campus buildings, including one dataset for cooling loads and one for ancillary information, as outlined below.

Campus Cooling Data

Our real energy dataset focuses on cooling load (given data availability) for 52 campus buildings, including dormitories, labs, offices, and academic buildings. The dataset begins on January 1, 2016, and ends on April 5, 2021, though pre-2020 data is selected due to COVID-19 irregularities. Several buildings are excluded due to data gaps, as described in the pre-processing section.

Ancillary Data

The historical hourly Boston weather data are downloaded from Oikolab¹, a weather API. The weather data include information such as dry-bulb temperature, dew point, relative humidity, solar radiation, rainfall, and snowfall, etc. These are potential features for time-series models as well.

Detailed campus building information is collected from the MIT Department of Facilities space accounting portal and the MIT Office of Sustainability DataPool, including construction year, building footprint, gross floor area, building type, and area breakdowns by room

¹<https://oikolab.com/>

use, etc. Specifically, the building use information could help deduct the usage schedule and thus the periodic pattern of the cooling load. Furthermore, we extract the engineering details from the metadata, such as construction material, air conditioning unit category (i.e. central chilled water, building chiller unit or local package units). This helps explain the variation among buildings. As is observed from Figures 9-1 and 9-2, campus building loads vary in both intensity and seasonality.




Figure 9-1: Ranges of hourly cooling loads for the 38 buildings on campus after removing buildings with outlier data²

9.2.2 Methods

Data Pre-processing

For the campus cooling load data, 21 of the 52 buildings are removed because of data gaps (with missing or unexpectedly many zero values). The remaining 31 buildings were optionally cleaned of outliers using a Hampel filter (24-hour window, 5 standard deviation threshold due to data variability), which generally corrected 2-5% of data points for the buildings without large gaps, as shown in Figure 9-2. Outlying points (in Figure 9-3) are distributed relatively evenly, with a slight weight towards times of peak cooling in the mid-afternoon hours and the month of June. Such a correction could help with model prediction without too much alteration, though analyses are also performed without the correction to test the raw data.

Model Features and Targets

We limit model features to the following for real data:

- Ambient dry-bulb temperature (T_{out}) [scaled]⁴

²Building 46 has a “whisker” range to ~ 2200 tons. Points beyond the whiskers are not necessarily outliers, but those simply outside the limit of $1.5 \times$ the inter-quartile range.

⁴A standard scaler that transforms the input data to a standard normal distribution to facilitate model training.

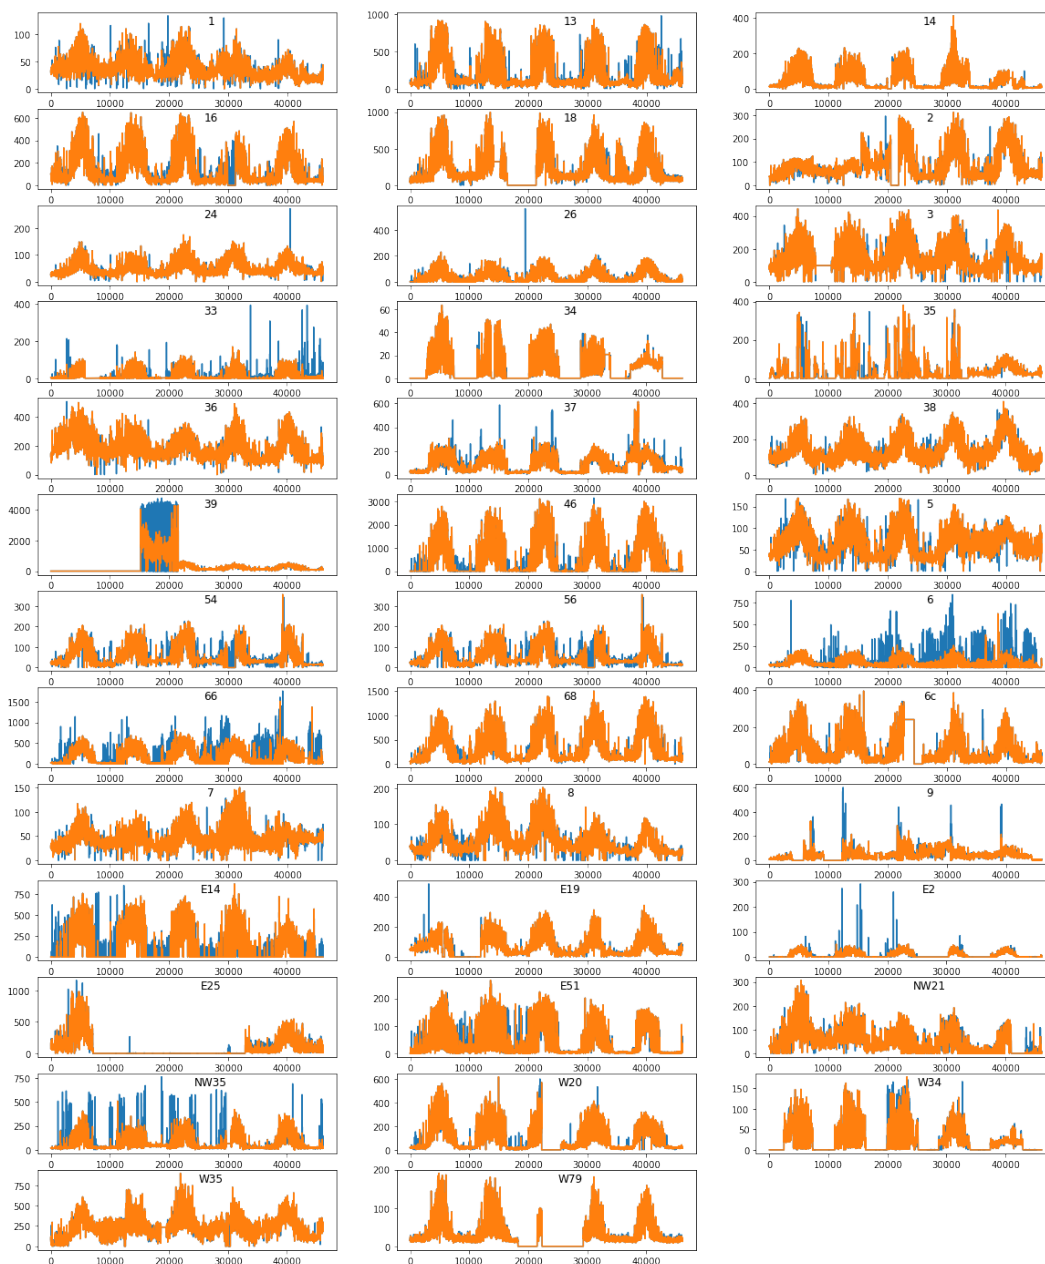


Figure 9-2: Cooling load time series for the 38 buildings, plotting before (blue) and after the Hampel correction (orange)³

³Buildings 18, 39, 6c, E25, and W79 are removed due to zero-cooling values at unexpected times (e.g. during the cooling season), as well as Buildings 6 and 35. Other buildings that have zero-cooling values outside of the cooling season are noted but kept, since this seasonality is expected.




Figure 9-3: Distribution of points corrected by the Hampel filter across all buildings, by month of year and hour of day.

- Temporal embedding with sine and cosine: time of day, day of week, month of year.

Due to occupant behavior, building schedule, and seasonality, building load demonstrates daily, weekly and seasonal trends. Consequently, we apply a sine and cosine embedding to all temporal information, as is discussed in [40]. For example, to represent "time of day", we use $[\sin(\pi(h)/12), \cos(\pi(h)/12)]$ to represent a 24 hour cyclic nature explicitly in the learning problem. In the case of the simulation dataset, we only encode time of day as a temporal feature.

We intentionally exclude T_{set} (room temperature setpoints) and dT_{set} (change of setpoints) because it is difficult to calculate an average T_{set} of a building containing hundreds of thermal zones⁵. Additionally, the dT_{set} for individual zones within a building is relatively constant during normal operation hours. The small perturbation and regular patterns of T_{set} indicates they are not the control factors of cooling load.

Time Series Models

We have explored three types of forecast model structures: a simple static multi-variate linear regression model, an ARIMA model, and a multi-layer perceptron (MLP)⁶. The 2017 hourly cooling dataset is used to train the model, and the performance is evaluated on the 2018 cooling load dataset.

1. Multi-variate Linear Regression (MLR) The target of the model is the HVAC load at time t . The features involve the ambient temperature, room temperature setpoint, and temporal information (time of day) at time t as well as previous timesteps. This model does not contain auto-regressive features.

⁵A thermal zone is defined as a space that shares the same temperature setpoint.

⁶MLP is a simple neural network structure.

2. Autoregressive with Exogenous Regressors (ARX(6, 1)) Similar to the method explored in Chapter 4, variations of the ARIMA model have been tested. Specifically, after comparing AR, ARIMA, ARX (AR with exogenous regressors), SARIMA (ARIMA with seasonality), and SARIMAX (ARIMA with seasonality and exogenous regressors), we find that including "Integration", "Moving Average" or "Seasonality" only marginally improves model performance. Considering the trade-off between model complexity and performance, we select ARX(6, 1) as the preferred method.

3. Multi-layer Perceptron (MLP) The neural network model in this chapter has the same structure as the one in Chapter 4. Overall, the inputs and output of the MLP follow the ARX(6, 1) model, and the only difference is that model fitting with MLP is non-linear whereas an ARX model is linear. It is an order-24 ARX model with order-25 exogenous inputs (including the current time step exogenous features). The forecasting horizon is 24 hours. We adopt this nonlinear structure rather than state-of-art time-series recurrent neural network (RNN) because RNN typically works best with high-dimensional data. In this specific learning problem where only a few features are considered, RNN performs much worse than a simple neural network structure.

9.2.3 Results


Performance

Model performance demonstrates the trade-off between simplicity and complexity, with neural network models outperforming AR(IMA)X, which out-performs linear regression. Root Mean Square Error (RMSE) values are compared in Table 9.1.


Table 9.1: Model performance (RMSE) on cooling hourly load prediction of Building 66

	MLR	ARX(6,1)	MLP
RMSE (Tons)	42.15	36.98	20.08

We have also noticed a difference in model performance in cooling vs. non-cooling seasons. As shown in Figure 9-4, the AR (6, 1) model performs much better in the summer cooling months than the winter non-cooling months, since the summer cooling load is much smoother and is driven mostly by ambient climate. More specifically, the RMSE for summer load is 7% compared to that of winter which is 50%. Given that summer cooling demand is also higher in general, this finding supports the use of such models for predicting and optimizing around these peak summer loads.



(a)



(b)

Figure 9-4: Hourly cooling load prediction of Building 66 with AR (6, 1) model for heating (a) and cooling (b) seasons

Sensitivity to Data Scale

In addition to data on Building 66, we also fit the ARX(6, 1) model on a group of buildings with continuous and positive cooling loads in order to examine whether prediction accuracy can change with aggregated data. Three building groups were selected:

- Building group 1: 5 buildings: 1, 5, 8, 24, 26
- Building group 2: 10 buildings: 1, 5, 8, 24, 26, 33, 36, 37, 38, 68
- Building group 3: 15 buildings: 1, 5, 8, 24, 26, 33, 36, 37, 38, 68, 66, 14, 16, 54, 56

The cooling data of all buildings in each of the three groups are summed before being fed into the model. Examining the scaled MSE and the percentage error ((RMSE / mean)

$\times 100\%$) of Building 66 and three building groups, we can see that overall, aggregated data yield better predictions (the error decreases each time more buildings are included). The scaled MSEs, unscaled RMSEs and percentage errors are summarized in Table 9.2.

Table 9.2: Summary of AR(6) model sensitivity to different prediction scales

Building	MSE^a	RMSE^b	Mean	% Error
Building 66	0.0492	36.98	170.66	21.67%
Building Group 1 (5 buildings)	0.0357	24.18	250.93	9.64%
Building Group 2 (10 buildings)	0.0188	76.90	957.38	8.03%
Building Group 3 (15 buildings)	0.0130	114.22	1461.59	7.82%

^aScaled error with features and targets scaled by a standard scaler.

^bUnscaled.

When a single building, such as Building 66 is examined, there are spikes of high cooling load periods that are not captured by the model. The spikes can result from operational abnormalities or special external conditions, such as building-specific use patterns or extreme outside temperatures. The building-specific spikes disappear as data from multiple buildings are aggregated together. On the other hand, the 15-building data show larger variation than the 5-building data, as the temperature-related peaks of each building are amplified by each other.

9.3 Campus Building Operations

In addition to the discussions on the methodological implications, the unique features of the MIT campus in terms of its building functionality, residents, and a semi-independent energy supply and sourcing system, deserve more attention when being integrated with the proposed control strategies. The campus energy supply and demand interaction is much more complicated than a single unit with an HVAC equipment, or a commercial and residential building with multiple units that consume electricity and natural gas directly from the grid. On the high level, MIT operates a tri-generation plant on campus that produces electricity and heat at low cost. Most campus buildings are on the campus district heating and cooling system. The following sections summarize the current campus conditions from the perspectives of demand (users) and supply (generation plant), present a preliminary feasibility analysis on implementing a data-driven optimal setpoint control strategy under the current conditions, and suggest next steps to prepare for the eventual launch of the initiative.

9.3.1 Demand Side

While it is tempting to invite more campus buildings to participate in setpoint adjustments, creating a larger potential for energy and GHG emission savings, it is critical to recognize

buildings have varied eligibility to be involved in this project due to hardware, software, and human constraints. We have included five major factors in Table 9.3 from the conversations with the Department of Facilities (DoF) and the campus leadership team to gauge the applicability to each campus building.

Table 9.3: Five factors for evaluating the eligibility of space-/building-level participation in the data-driven setpoint control project

Decision factors	Campus conditions and criteria	Notes
BMS availability	(1) Room-level control (2) Compatibility with StruxtureWare	Some other BMS platforms will be migrated to StruxtureWare by 2025 as planned.
HVAC equipment condition	(1) Age < 25 years (2) Type: VAV, Fan coil unit (FCU), chilled beam, FTR, and some system combinations	CAV dual-duct and valance units are excluded.
Space type	(1) Eligible: conference rooms, offices, and classrooms (2) Waitlist: dormitories and labs	-
Occupants' response	(1) Occupant's current level satisfaction with building HVAC operations	Buildings 12 and E37 are excluded.
Others	(1) Limitations of the data-driven approach (2) Feasibility of integration of other control algorithms with the data-driven one	-

I. Availability of Building Management System (BMS)

All the setpoint control methods require buildings to have room-level BMS control; that is, all the room conditions and setpoint controls are accessible through the BMS software. Currently, MIT campus buildings have multiple BMS platforms, such as StruxtureWare (the one with Building 66 experiment), Continuum, ALC, and Siemens⁷. At this moment, buildings with StruxtureWare and some with Continuum and ALC are able to provide room-level monitoring, which enables approximately 40 buildings to participate⁸.

⁷The detailed documentation of BMS for each building is included in the Appendix C.

⁸MIT DoF plans to migrate Continuum and ALC platforms to StruxtureWare by 2025.

II. Conditions of HVAC Equipment

Depending on building construction time, the installed HVAC equipment may not support the setpoint adjustment framework. The two interconnected decision factors are (1) age of equipment and (2) type of equipment. The older buildings (i.e. constructed 100 years ago) are installed with HVAC equipment that cannot be accurately controlled on the room-level. In that regard, a 25-year equipment service age is the preliminary threshold recommended.

Most aged buildings on campus installed dual-duct constant air volume systems (CAV) that supply cold and hot air (conditioned by central AHU to fixed temperatures) simultaneously in the two decks and mix cold and hot air to get to the desired supply air temperature. The setpoint changes on the room level can then lead to changes in building-level hot deck and cold deck temperature setpoints to reduce energy consumption. However, most buildings with dual-duct CAV systems do not support setpoint control due to their outdated control hardware and software instrumentation. Therefore, they are not a prime target in the future research scope.

Another system type discouraged in the project is the valance unit in the new graduate students' apartment⁹. Introduced only recently to campus, valance units are not very welcomed by residents due to their low air-circulation conditioning feature. While minimizing airflow allows energy savings, the system creates a stagnant feeling and reduces the perceived cooling by convective flow passing human skin. Based on occupants' feedback, the campus no longer considers valance units in new campus construction.

Other system types, such as VAV, FCU, chilled beam, and FTR are eligible, but worth further investigation through reading control specifications and building mechanical drawings, and performing commissioning experiments. For buildings that deploy combined system types, such as chilled beam and FTR with VAV in Building 66, it is especially crucial to understand the control logic and coordination among all systems. We highly recommend prioritizing single-system space experimentation to multi-system buildings that require far more attention and commissioning.

III. Type of Space

Space with various functionalities lead to different requirements on thermal comfort and indoor environmental quality (IEQ). Lab space with the most stringent IEQ standards needs a comprehensive evaluation of the consequences associated with setpoint and airflow changes. Nonetheless, labs, being the most energy intensive space type, have the largest potential in energy and GHG savings, if control methods are carefully designed and executed. Dormitory, where students spend most of their time studying and resting, is also worth further discussions. While algorithms may achieve significant savings during the unoccupied hours

⁹Graduate Tower at Site 4, Building E37

in dorms, mis-prediction or optimization errors may lead to significant impacts on thermal comfort and impose burden to building managers and administrators. On the other hand, conference rooms, offices, and classrooms that have regular schedules and more predictable loads are more suitable for setpoint adjustments. In the meantime, the gatherings in classrooms and conference rooms are not expected to last long, so setpoint adjustments would have limited influence on space users. In case thermal discomfort occurs, space users can immediately file a thermal comfort survey posted in the room.

IV. Response and Perception of Occupants

The perception of building occupants also plays an essential role that influences the success of implementation in practice. Based on current knowledge, we would recommend avoiding instrumentation in buildings whose occupants have already demonstrated dissatisfaction toward the indoor environment, such as Building 12 (the Nano building), and Building E37 (Site 4 graduate dorm). To consider occupants' response, we suggest launching a campus-wide survey to gauge how much acceptance they have towards setpoint adjustments, and what the research team and DoF could do to resolve their concerns and minimize implementation challenges.

V. Other Considerations

Prior to the data-driven setpoint control project, DoF initiated a series of HVAC setpoint control upgrades to fulfill the goal set by the MIT Climate Action Plan. Specifically, a majority of campus buildings have switched from single-setpoint control to dual-setpoint control (with separate cooling and heating setpoints that allows for a larger HVAC operation deadband). In the meantime, pilot occupancy-based setpoint control experiments have been conducted in selected buildings, which demonstrate non-negligible energy savings. As is elaborated in Chapters 4 and 6, the data-driven algorithms have been adjusted and tailored to HVAC operation in heating and cooling seasons during the onsite experiment but may not be well integrated with the dual-setpoint control scheme during the shoulder season. In addition, we would like to explore how our data-driven approach can be incorporated with the occupancy-based control to achieve further savings. We envision that all the methods tested and validated are not independent of each other; MIT personnel should not hesitate to create synergistic effects by integrating them appropriately, if there are ways to do so.

9.3.2 Supply Side

Besides considerations from the energy user side, the MIT campus deploys a complex energy supply system, without investigating into which the true energy and GHG reduction goal cannot be achieved. MIT operates a tri-generation plant, producing electricity and heat on

natural gas. In the meantime, it sources electricity from the grid when grid electricity price is low. This has several implications and prompts us to discuss the following questions.

I. How can the cost-based supply side Optimization fulfill a carbon reduction objective?

The current energy procurement is based on cost optimization, whereas the proposed control methods leverage real-time grid carbon intensity for optimization. A misalignment in optimization objective would lead to a sub-optimal GHG saving outcome. Therefore, MIT is advised to re-evaluate and re-design the energy procurement decision model to an either carbon-based one, or multi-objective optimization framework that considers both factors. In addition, the existing load forecast model used for electricity procurement cost minimization will not reflect load perturbation when setpoint adjustments are implemented. To support load perturbation mechanisms and make an informed procurement decision, it is necessary to train a new or upgrade the existing load forecast model with additional features (i.e. setpoint adjustments) from building operations side.

II. How to allocate and estimate the carbon emissions associated with different energy streams generated from a tri-generation plant?

The energy efficiency of a tri-generation plant comes from the fact that the plant makes use of the engine heat from electricity generation to produce steam or hot water for most campus buildings. In winter, hot water can be supplied to heating systems to provide free heating energy. In summer, some chillers on campus run on steam to provide cooling. Currently, MIT attributes all carbon emissions from a tri-generation plant to electricity and considers the bi-products steam and hot water as carbon-free. This calculation framework needs to be updated to fairly assign a carbon value to steam and hot water, reflecting the carbon savings associated with reduced cooling and heating loads.

In addition, a decrease in electricity consumption associated with cooling load decrease may lead to a loss of free heating energy available on campus, and therefore may require additional energy consumption to make up for heating load requirements. To fully understand this topic, we need to look carefully at the plant-level operation, the performance metrics of the plant, and the consumption rate of electricity, hot water and other related energy streams on campus. This ensures we maintain the synergies of a tri-generation plant and make them suitable to the new operation framework.

III. How can the control method be updated when campus heating or cooling loads are not fulfilled by electricity?

The prototype optimization algorithm executed in the onsite experiments makes optimal control decisions based on the grid carbon intensity. In practice, the HVAC systems in campus buildings involve various types of energy sources that cannot be easily related to grid carbon intensity. First, the electricity used for cooling comes largely from burning natural gas, with others from direct procurement from the grid. Second, some campus chillers run on electricity, while others on steam produced by burning gas. In winter, heating supply on campus is entirely from natural gas. To implement the optimization framework, further adjustments need to be made to reflect the complexity of the energy supply system on MIT campus. Specifically, we need to figure out the objective of the optimization when grid carbon intensity is irrelevant. This could be minimizing total heating loads, peak heating load demand to maintain plant efficiency, or other factors that may fit the goal of carbon emission reduction initiative better.

Chapter 10

Conclusion and Future Work

This thesis explores the feasibility of three optimization-based control methods in simulation testbeds for direct load control for a large utility company and MIT campus to regulate load demand to achieve peak load reduction and GHG emissions savings. In addition, eight onsite experiments are conducted to verify the effectiveness of one of the proposed control algorithms.

10.1 Conclusion

The main contributions of the thesis include:

1. This thesis proposes and validates three optimization based control algorithms, namely the three-step "predict-perturb-optimize" approach, the two-step "predict-optimize" approach, and the model predictive control (MPC) method. Simulation results demonstrate the capability of the algorithm in reducing peak load demand by 20% to 40% and minimizing GHG emissions by 30% to 50%, while fulfilling thermal comfort and indoor air quality requirements.
2. This thesis manages to execute onsite experiments to further validate the methods in practice, which showcases a 30% to 60% load savings depending on the system types. Experimental validation is uncommon and valuable in this field because hardware and software setup can be expensive and laborious, and data collection, processing, and interpretation could involve numerous challenges.
3. On the simulation side, the simulation testbed from Chapter 3 provides a useful and convenient numerical based testbed for future research efforts on aggregate load demand control. Although this is a 1RIC model, it integrates important parameters, such as solar heat gain, HVAC COP, and heat transfer between zones to allow for a more realistic but computationally efficient simulation model. It also establishes a

pipeline where users can parametrically define their simulation scenarios, perturbation magnitudes, and optimization targets in a sequential framework for a wide variety of research purposes.

4. The simulation based MPC approach explored in Chapter 5 addresses a multi-objective optimization problem that considers not only the commonly discussed targets, such as thermal comfort, energy cost, or GHG emissions, but also indoor air quality. This MPC model incorporates both thermal dynamical simulation and a CO₂ mass balance model combined with convex relaxation techniques, which allows for dynamical adjustment of room setpoints and room CO₂ concentration level to meet requirements on peak load and GHG emissions.
5. On the experiment side, the thesis discusses hardware, software, and implementation issues, critiques the experiment design, and finally provides multiple feasible and practical load saving validation methods. This includes valuable lessons learned from onsite experiments, especially in education institutions where energy and supply systems tend to be more intricate than typical residential and commercial community. Future researchers can refer to Chapters 6 to 9 to gain insights from the experiment protocols, methodology design, and result analysis to make onsite experiments less challenging and more productive.

10.2 Future Work

Notwithstanding the merit and contributions, future work can be investigated in the following directions as an extension of the thesis:

1. On the simulation side, the thesis mainly adopts convex and linear optimization as the setpoint control optimizer. As is noted in Chapters 3 and 4, a non-linear (non-convex) optimizer could out-perform a linear one when the embedded surrogate non-linear model achieves a higher prediction accuracy than a linear surrogate model. Researchers are encouraged to integrate non-linear surrogate models with non-linear optimization algorithms to seek larger savings on energy consumption and GHG emissions.
2. The convex relaxation method used in Chapter 5 to deal with the non-convexity of room CO₂ dynamics can be improved by using Taylor approximation method. This ensures a more stable prediction and optimization and allows for a larger optimization horizon to further promote model performance.
3. On the experiment side, sensor and meter calibration should obtain constant attention as issues regarding flowmeter sensitivity level and reading differences across two test

rooms are unresolved. It is suggested that the entire hardware and software should be maintained robust to provide consistent readings for load saving analysis.

4. More experiments are needed to collect data points for validating the control and experiment methods and obtaining reliable room thermal properties such as the UA value described in Chapter 8.
5. MIT and future research teams are encouraged to establish a data pipeline that automatically collect different data streams (i.e. weather forecast, real-time occupancy data, thermal loads, and grid carbon intensity) from various sources to streamline the experiment setup and make the process more efficient.
6. All the methods and experiment verifications serve as proof of concept and need to be investigated for large-scale implementation (such as on the entire MIT campus). Choosing the suitable method for aggregate implementation is critical as it involves in-depth analysis of the site conditions and restrictions. For example, the discussions in Section 9.3 involve a comprehensive list of topics for MIT to think through for materializing the savings.

Appendix A

Experiment Protocol and Setup

A.1 Typical Weekend Test Setup¹

Below is an example Experiment Checklist for Oct 23-24 weekend test, with control over room condition, including room access, lighting, blind position, and projector following the checklist below.

Checklist

1. Signages up on the connecting doors in Rooms 154, 156, and 160, front doors in Room 154 and 160, and near light switches²
2. Reserve the space with Registrar Office
3. Confirm with custodians to lock doors for weekends on Friday evening
4. Check room conditions (Friday evening, Saturday and Sunday mornings)
 - (a) Lights off
 - (b) All blinds down
 - (c) Front and connecting doors closed
 - (d) Projectors off
5. Signs off before 9am Monday
6. Saturday: baseline test
7. Sunday: optimization + simulation test
8. Result analysis:
 - (a) VAV load
 - (b) Chilled beam load
 - (c) Radiator load

For other experiments with less controlled condition, we only put signages in the rooms requesting students to follow the instructions when using the rooms. We no longer reserve

¹Depending on the experiment date, protocol may vary.

²All signages used in this and other tests are attached in Figure A.1

the space or lock out rooms.






- 
EXPERIMENT USE
Please ***DO NOT*** leave the door open
Thank you! 😊
For any questions, please contact
Yuan Cai @ 512-757-1898 or yuancai@mit.edu
- (a) The sign on the connecting doors
- 
EXPERIMENT USE
Please follow the ***INSTRUCTIONS*** in the room
when self/ group studying
Thank you for cooperation! 😊
For any questions, please contact
Yuan Cai @ 512-757-1898 or yuancai@mit.edu
- (b) The sign on the door facing hallway
- 
EXPERIMENT USE
Please ***DO NOT*** touch or remove the **sensors** in
the room
Thank you! 😊
For any questions, please contact
Yuan Cai @ 512-757-1898 or yuancai@mit.edu
- (c) The sign put on the front desk to remind space users of the illuminance sensors in the rooms
- 
EXPERIMENT USE
Please ***DO NOT*** enter the room
Please find other space for group study
Thank you for cooperation! 😊
For any questions, please contact
Yuan Cai @ 512-757-1898 or yuancai@mit.edu
- (d) The sign on the hallway door requesting students not to use the space during weekends
- 
EXPERIMENT USE
Please ***DO NOT*** change the **blind** position, turn
on the **lights** or **projector**
If you do need to use them, please (1) **FULLY**
turn up the two blinds; (2) **turn off lights**; (3)
turn off projectors before leaving the room.
- Thank you for cooperation! 😊
- (e) The sign for room usage instructions near the light switches

Figure A-1: Signage examples

A.2 Illuminance and Temperature Sensors (HOBO) Placement

The experiments on Nov. 20-21 and Nov. 25-57 deployed HOBO (illuminance and temperature) sensors that monitored the solar heat gain through windows and lighting conditions. In total six HOBO sensors were placed in the two rooms, two on the windowsills and one on the front desk, as is illustrated in the Figures A-2a and A-2b. The two sensors lay horizontally on the windowsills between the curtain blinds and the glass panels, so they were not set up to detect the changes in the blind positions.

Room 154



Senor
Serial #

- Notes:
1. Each room has 3 sensors, 2 on the window sills, and 1 on the desk in front of the white/black board.
 2. The 2 sensors on the window sills are between the window panels and the curtain blinds.

Unlabeled

(a) [Room 154] Sensor layout

Room 160



Senor
Serial #

- Notes:
1. Each room has 3 sensors, 2 on the window sills, and 1 on the desk in front of the white/black board.
 2. The 2 sensors on the window sills are between the window panels and the curtain blinds.

TS9

(b) [Room 160] Sensor layout

Figure A-2: HOBO (illuminance and temperature) sensors placement location in Room 154 (a) and Room 160 (b)

Appendix B

Experiment Detailed Documentation and Results

Detailed load excel spreadsheet calculations can be download from the GitHub Repo: <https://github.com/YuanC233/Thesis-setpoint-optimization.git>

B.1 Experiment Documentation

B.1.1 Sep. 8 (Cooling Season)

Purpose

To test the feasibility of setpoint control in the two rooms (Room 154 as optimal control and Room 160 as baseline) to minimize load and GHG emissions

Test Setup

1. Adjusted heating and cooling temperature setpoint simultaneously, 0.5 °F of gap (the system does not require a minimal setpoint gap between heating and cooling)
2. Baseline experiment: Room 160; optimized control: Room 154
3. Room 160: 24-hour cooling setpoint at 69.8 °F; heating setpoint at 69.3 °F
4. Room 154: when room is occupied, cooling setpoint should be between 68.0 °F and 69.8 °F, otherwise, between 64.4 °F and 80.6 °F
5. Little instructions on room usage and control over room occupancy, lighting, or other conditions

Results

1. Identified system control and operational issues such as VAV operating as CAV, chilled beam limited cooling capacity, and inability to track room setpoints

Notes

1. This experiment only provided VAV related room loads because the BTU meters for chilled beam and FTR were not installed or calibrated yet

B.1.2 Oct. 10 – Oct. 11 (Cooling Season)

Purpose

To verify if Room 160 could serve as a baseline reference for Room 154, the critical hypothesis for running a parallel test discussed in Chapter 6

Test Setup

1. Both rooms operated under a 3 °F deadband with cooling setpoint at 72.5 °F and heating setpoint at 69.5 °F
2. The minimal air flow provided by the VAV system dropped from 450 cfm down to 250 cfm; the maximal air flow raised from 450 cfm to 700 cfm
3. Room usage instructions were provided

Results

1. Based on VAV load analysis, the two room loads demonstrated similar patterns except during the noon time when solar exposure caused load disturbances
2. Observed the abnormality that room loads were driven by the air volume supplied by VAV from data collected in Rooms 154, 156, and 160

Notes

1. This experiment still did not provide BTU meter data on chilled beam or FTR

B.1.3 Oct. 16 – Oct. 17 (Cooling Season)

Purpose

To continue verification of similarity of load patterns of Room 154 and Room 160, with more rigorous room condition control

Test Setup

1. (Same as the experiment on Oct. 10 – Oct. 11) Both rooms operated under a 3 °F deadband with cooling setpoint at 72.5 °F and heating setpoint at 69.5 °F
2. Students were advised not to enter the rooms during the testing period; room conditions were monitored everyday

Results

1. Even with a stricter room condition control, the VAV load Pearson correlation was still around 70%, which indicates that the impact of solar loads is non-negligible

Notes

1. This experiment still did not provide BTU meter data on chilled beam or FTR

B.1.4 Oct. 23 – Oct. 24 (Shoulder/Heating Season)

Purpose

The first experiment that adopted a 3 °F deadband on the optimal control day in a shoulder/heating season

Test Setup

1. Oct. 23 (baseline day): Room 154 and Room 160 both deployed a constant cooling setpoint at 69 °F and a constant heating setpoint at 72 °F
2. Oct. 24 (optimal control day): Room 154 deployed an optimized heating and cooling setpoint schedule, while Room 160 remained in baseline mode
3. The 2-step optimization framework only output the heating setpoint schedule; because of the constant 3 °F deadband, the cooling setpoint schedule also changed with the heating setpoints; however, cooling setpoint was not a result of the optimization

Results

1. Adjusting heating and cooling setpoints simultaneously introduced larger energy consumption (VAV + chilled beam + FTR) on the optimal control day, which is opposite to our experiment goal

B.1.5 Nov. 20 – Nov. 21 (Heating Season)

Purpose

To test a varying setpoint deadband with heating setpoint defined by optimization algorithm and cooling setpoint constant

Test Setup

1. Followed a typical two-day test schedule
2. Baseline heating setpoint: 69 °F (constant), cooling setpoint: 72 °F (constant)

3. Optimal control heating setpoint: optimization output (varying), cooling setpoint: 76 °F (constant)
4. Students were informed to keep classrooms in good experimental condition with signs on
5. HOBO (illuminance and temperature) sensors installed to track room lighting and solar heat gains

Results

1. The varying setpoint deadband strategy turned out to be successful; energy savings were identified during the night period when heating setpoints were adjusted downwards

Notes

1. Due to implementation issues, the optimal control experiment only started after 5pm, so results were analyzed based on the 7-hour period

B.1.6 Nov. 25 – Nov. 27 (Heating Season)

Purpose

To validate energy savings with a varying setpoint deadband control strategy

Test Setup

1. Followed a typical two-day test schedule except that the second optimal control day got replicated on the third day to collect more data points
2. Baseline heating setpoint: 69 °F (constant), cooling setpoint: 72 °F (constant)
3. Optimal control heating setpoint: optimization output (varying), cooling setpoint: 76 °F (constant)
4. Students were informed to keep classrooms in good experimental condition with signs on
5. HOBO (illuminance and temperature) sensors installed to track room lighting and solar heat gains

Results

1. Energy savings were identified, but a more systematic evaluation method was needed because it was found out that the heating loads from FTR of the two rooms were not comparable (different by a factor of 5)

2. A ballpark estimation concluded an approximately 20% - 50% savings on the GHG emissions and thermal load; however, due to the potential errors in the sensor measurements, the findings were mixed and could not be generalized with confidence

Notes

1. Sensor recalibration and reinstallation was conducted after the experiment
2. Because the optimization is entirely conducted on the simulated data (without knowing the historical room setpoints and temperature in the classrooms in reality), the back-to-back optimal control tests would introduce inconsistent setpoint control strategy associated with higher thermal loads. To avoid this situation, it is recommended that a break is designed between the optimal control tests, or thermal loads are only calculated based on a 23-hour period to exclude the transition period (usually the first hour of the day)

B.1.7 Dec. 26 – Dec. 30 (Heating Season)

Purpose

To validate energy savings with a rule-based benchmark

Test Setup

1. Followed a typical three-day test schedule (baseline – rule-based control – optimal control)
2. Rule-based heating setpoint (in both rooms): 64 °F when unoccupied and 69F when occupied, cooling setpoint: 76 °F when unoccupied and 72 °F when occupied
3. Baseline heating setpoint: 69 °F (constant), cooling setpoint: 72 °F (constant)
4. Optimal control heating setpoint: optimization output (varying), cooling setpoint: 76 °F (constant)
5. Students were informed to keep classrooms in good experimental condition with signs on

Results

1. After sensor recalibration, the results were more interpretable and generalizable. Optimal control day showed a 30% - 60% reduction in room loads, and rule-based control showed a 70% reduction compared to the baseline.

Notes

1. Test schedule: Dec. 26 – baseline, Dec. 27 – optimal control with implementation errors (thus excluded from result analysis), Dec. 28 – test break, Dec. 29 – rule-based control, Dec. 30 – optimal control

B.1.8 Jan. 1 – Jan. 3, 2022 (Heating Season)

Purpose

To repeat experiment series on Dec. 26 – Dec. 30

Test Setup

1. Same as the setup in the last experiment

Results

1. Results were similar to the ones from the last experiment, indicating this conclusion is reproducible

Notes

1. Test schedule: Jan. 1 – rule-based control, Jan. 2 – optimal control, Jan. 3 – baseline
2. Unresolved but minor sensor issues still existed, but did not impact experiment results

B.2 Room Load Sample Calculation (Chapter 8)

B.2.1 FTR and Chilled Beams

Calculation of heating and cooling loads from FTR and chilled beams are straightforward because those values are pre-calculated in the BTU meter database (downloaded from EcoStructure). All the data are recorded and calculated on a 5-minute interval.

Use the heating load of FTR (Q_{FTR}) in Room 154 as an example.

The values in the column “Rm154_HwBTU_Calculated Energy_Tr_AITr (Btu)” are cumulative thermal loads from FTR. You can get hourly load by adopting a center differencing method.

Table B.1: Sample data

Time	Rm154_HwBTU_... (Btu)
1/3/2022 12:30 AM	37216116.00
1/3/2022 12:35 AM	37217064.00
1/3/2022 12:40 AM	37217376.00

The hourly thermal load at 12:35 AM can be calculated by:

$$\begin{aligned} (\text{Load at 12:40 AM} - \text{Load at 12:30 AM}) \times 6 &= (37217376.00 - 37216116.00) \times 6 \\ &= 7560 \text{ BTU/h} \end{aligned}$$

The same method can be applied to cooling load from chilled beams (Q_{CB}).

B.2.2 VAV

The VAV loads are not readily computed so they need to be manually calculated using formula:

$$Q = c_{air} \dot{m}_{air} \Delta T \Delta t = c_{air} \rho_{air} \dot{V}_{air} (T_{room} - T_{supply}) \times (t_e - t_s) \quad (\text{B.1})$$

Where \dot{m}_{air} , \dot{V}_{air} refer to mass flowrate and volumetric flowrate of air supplied by VAV, T_{room} is room temperature, T_{supply} is supply air temperature of VAV, t_e and t_s are the end and start time of interest, c_{air} and ρ_{air} are specific heat capacity and density of air.

Cooling Season

In the cooling season, when the outdoor air supplied by VAV provides cooling and maintains indoor CO₂ level, the cooling load can be directly calculated using Equation (B.1).

For example:

$$\begin{aligned} \text{Rm154_VAV.SaFl (CFM)} &= 252.0 \\ \text{Rm154_VAV.SaTmp (}^\circ\text{F)} &= 55.0 \\ \text{Rm154_Chb/FTR.RmTmp (}^\circ\text{F)} &= 72.0 \end{aligned}$$

$$\begin{aligned} Q_{VAV} &= c_{air} \rho_{air} \dot{V}_{air} (T_{room} - T_{supply}) \times (t_e - t_s) \\ &= 0.240 \text{ BTU/}^\circ\text{F-lb} \times 0.0754 \text{ lb/ft}^3 \times 252.0 \text{ ft}^3/\text{min} * (72.0 - 55.0) \text{ }^\circ\text{F} \times 60 \text{ min} \\ &= 4651 \text{ BTU/h} \end{aligned}$$

Heating Season

For the heating season, the computation is more complicated because some heating or cooling loads provided by VAV are not intended to track setpoints, but to bring fresh air into the space. To separate the two types of flows, a filtering algorithm is applied with the following logic. Note that the filtering logic is not able to perfectly separate the flows because the collected data involve uncertainties and trends that are not easily generalized with simple logic statements.

To filter out the effective or active cooling flowrates, two conditions need to be fulfilled simultaneously. (a) the room temperature is higher than the cooling setpoint, and (b) the

supply air flowrate is higher than the minimal flowrate. Because of the deadband effect (the cooling system will not be triggered immediately after the room temperature exceeds the cooling setpoint), the actual cooling setpoint boundary defined in the filtering algorithm may be +0.5 °F or +1 °F higher than the one in the system. Likewise, while the minimal flowrate is kept at 250 cfm, the actual damper control may bring 270 cfm as the observed minimal flowrate. So the flowrate cutoff should be set at any number that is reasonable (275 cfm is used in this analysis).

To filter out the effective heating flowrates, three factors need to be satisfied: (a) supply air temperature should be above the room temperature, (b) room temperature should be lower than the heating setpoint, and (c) the VAV reheat valve position should be open. Again, the three cutoffs may require certain tolerance considering the measurement errors and deadband effect.

The effective cooling and heating flowrates then can be fed into Equation (B.1) for the respective VAV load calculation.

B.3 Predictive Baseline Load Calculation and Load Saving Estimation (Chapter 8)

The load saving analysis for Room 154 uses a predictive baseline with load information from Room 160. Given the strong correlation between Room 160's and Room 154' heating loads (observed in Figures 8-11), we use the slopes to scale Room 160's baseline load on the optimal control day of Room 154 to define the baseline load for Room 154.

To use FTR heating load as an example.

$$Q_{FTR_154} = 4.69 \times Q_{FTR_160} - 17403 \quad (\text{B.2})$$

On Dec. 30, Room 154 was implemented an optimal control test, while Room 160 remained as the baseline. Room 160's baseline Q_{FTR} is 15364 BTU. With Equation B.2, Room 154's baseline Q_{FTR} can be predicted as 54654 BTU. As the optimal control load is 36128 BTU, the total savings on FTR load is 34%.

Appendix C

Documentation of External Sources

C.1 Psychrometric Optimization for Supply Air Dry Bulb Set-point for Buildings 18, E25, and 66




Figure C-1: Psychrometric optimization for supply air dry bulb setpoint for Buildings 18, E25, and 66

C.2 Sequence of Operations for HVAC Equipment (from DoF and Schneider Electric)

Typical for Classrooms Floor 1

- General:
 1. This control scheme consists of an Active Chilled Beam with zone pump fed by a VAV box with reheat coil, fin tube radiation, RH control, CO2 control.
 2. The BMS will monitor room temperature using a standard smart room sensor with display.
 3. The room sensor shall provide room setpoint override capability, and a room occupancy override of 4 hours
 4. The chilled beam piping will be equipped with a pipe temperature sensor located on the discharge side of the zone pump. In general, the CHW valve will modulate to maintain space temperature setpoint. The low limit for mixed water temperature setpoint shall be determined by the outside air reset schedule below but always remain 2 degrees above space dewpoint.

Outside Air Dew Point	Mixed Water Set Point Low Limit
55 Degrees F	59 Degrees F
52 Degrees F	55 Degrees F

- 5. The VAV box shall normally control to a fixed set point based on scheduled CFM with space CO2 and space dewpoint having the ability to override normal CFM set points with space dewpoint taking priority.
- VAV Minimum CFM Reset:
 1. When no occupants are detected and the space is deadband mode, the VAV damper shall be commanded closed and the flow set point set to the flow to prevent any alarms.
 2. On detection of an occupant, the VAV minimum CFM set point will be set to the flow resulting in a velocity pressure of 0.025
 3. The minimum CFM set point shall reset based on the number of occupants by 50 CFM per occupant until the minimum CFM set point reaches the design minimum CFM set point.
- Occupied Mode:
 1. Occupied mode will be determined via BMS Schedule (adj.).
- Temperature Control:
 1. Occupied/Daytime/Daytime Standby
 - a. During occupied mode the space will incorporate a deadband range, typically 3 degrees, which establishes the space heating and cooling setpoints. Heating setpoint is 71 and cooling is 74. (3 deg deadband)
 - b. Deadband operation. During the deadband temperature range the zone pump will be shut off once the CHW valve is fully closed and space temperature setpoint has been satisfied for 15 minutes (adj.). The CHW valve, VAV RH valve, FTR valve, and VAV damper shall be closed.
 - c. On a rise in room temperature above occupied cooling setpoint, the CFM set point shall reset to minimum and the zone pump shall turn on. The CHW valve shall modulate to maintain mixed water temperature setpoint until space temperature setpoint is met. If space dewpoint comes within 2 degrees of mixed water temperature, the mixed water temperature setpoint shall be increased in 1-degree increments to remain above the dewpoint from 59 degrees to 74 degrees.

Figure C-2: Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160


- d. If space dewpoint comes within 1 degree of mixed water temperature, the zone pump shall be commanded off, CHW valve shall close and the VAV box shall open 100% until space dewpoint falls 3 degrees below mixed water temperature setpoint degrees prior to starting the zone pump and opening the chilled water valve again.
 - e. If CHW valve is full open and space temperature continues to rise CO2 control shall be overridden and the VAV box shall modulate open.
 - f. As room temperature begins to drop towards occupied cooling setpoint, the VAV box shall be released from temperature control and the CHW valve shall modulate closed in the reverse order.
 - g. On a drop in room temperature below occupied heating setpoint the VAV box shall modulate its damper to minimum position required to maintain space CO2 setpoint, CHW valve shall modulate closed. If the CHW valve is fully closed and the space temperature continues to drop, the RH valve shall modulate open to maintain discharge air temperature setpoint, which shall be reset by the space temp. If further heating is needed, the FTR valve shall modulate open.
 - h. When outdoor air temperature is below 50 degrees the RH valve and FTR valve shall modulate open in unison
 - i. As room temperature begins to rise towards occupied heating setpoint, the FTR valve and RH valve shall modulate closed in the reverse order.
2. Unoccupied:
- a. When room is unoccupied, the room temperature cooling setpoint shall be reset to 80 Deg F. (adj.) and the room temperature heating setpoint shall be reset to 65 Deg F. (adj.). The zone pump shall be commanded off and the CHW, HW, RH, FTR valves shall be closed.
 - b. On a rise in room temperature above unoccupied cooling setpoint or rise in space dewpoint above 59 degrees, the VAV shall modulate 100% open.
 - c. Once the room temperature has reached 3 Deg F. below unoccupied cooling setpoint, the VAV box shall close.
 - d. On a drop of room temperature below unoccupied heating setpoint, the VAV box shall modulate open 100%, RH and FTR valves shall open 100%.
 - e. Once the room temperature has reached 3 Deg F. above unoccupied heating setpoint, the VAV box shall modulate to minimum position, RH and FTR valves shall close.
- CO2 Control
 - 1. The BMS will monitor space CO2 via wall mounted sensor.
 - 2. If the CO2 level rises above 850 PPM, the VAV damper shall modulate towards maximum to maintain a 450 PPM until the VAV is down to minimum CFM setpoint.
 - Sensor Override
 - 1. Each space will be equipped with a Smart sensor and will operate as outlined in the instruction sheet below.

Figure C-3: Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160 (continued)

Bldg 66 1st Floor Classroom Thermostat Operation

The following user adjustments are available

*** Temp setpoint adjustment * Unoccupied override**



Temperature Setpoint Change

1) Press ▲ or ▼ to increase/decrease temp, then ENTER
* See Heating & Cooling setpoints limits below

Unoccupied Override

1) Press button
* Unoccupied Override will last 4 hours, at which time it will return to the Unoccupied mode.

Setpoints

	Min Setpoint	Max Setpoint
Occupied Cooling	70.5	75.5
Occupied Heating	67.5	72.5
Unoccupied Cooling	78	
Unoccupied heating	64	

• Alarms

1. The following alarms will be reported to the BMS
 - a. Room temperature deviates by +/- 10 Deg F. from room temperature setpoint for 10 minutes.
 - b. Space dewpoint rises within 1 Deg F of chilled water temperature.
 - c. A critical alarm shall be generated if space dewpoint rises 1 degree over chilled water temperature.
 - d. Zone Pump commanded ON but status shows OFF.
 - e. Space CO2 rises above 1200 PPM.

Figure C-4: Control sequence documentation for a combine VAV, chilled beam, and FTR in Room 154 and Room 160 (continued)

C.3 Sensor Specifications and Documentations

C.3.1 Flowmeter Sensor: U1000MKII-FM Fixed Clamp-on, Ultrasonic Flow Meter

Website: <https://micronicsflowmeters.com/product/u1000mkii-fm-fixed-clamp-on-ultrasonic-flow-meter/>

The New ULTRAFLO U1000MKII-FM is a Fixed, clamp-on, pipe mounted, ultrasonic flow meter, which delivers significant installation savings and non-invasive, dry servicing benefits for pipe sizes ranging from 22 mm – 180 mm (6”) OD.

Technical Information

Measurement Technique: Ultrasonic, cross-correlation transit time method for flow measurement.
Turn Down Ratio: 100:1
Accuracy: $\pm 1\%$ – 3% of flow reading for $>0.3\text{m/s}$ [$>1\text{ft/s}$].
Flow Velocity Range: 0.1m/s – 10m/s [0.3ft/s – 32ft/s].
Pipe Range: Available in 2 options. 22mm – 115mm and 125mm – 180mm OD. Note Pipe size dependent on pipe material and internal diameter.
Pipe Material: Steel, Stainless Steel, Plastic and Copper
Water Temp Range: 0°C – 85°C [32°F – 185°F].
Pulse Output: Pulse or Frequency. Pulse for volume flow and alarms. Frequency for flow rate. The pulse output can be configured as loss of signal or low flow alarm. Opto-isolated MOSFET volt free contact [NO/NC].
4-20mA Output: Optional 4-20mA flow proportional output. Resolution 0.1% of full scale. Maximum load 620 Ω .
Modbus Communication: Pipe Mounted:- Optional Modbus RTU slave, RS485 serial link hardware layer. Connection cable 1M.
M-Bus Communication: Optional M-Bus
External Power Supply: 12V-24V $\pm 10\%$ AC/DC at 7 watts. Optional plug in 12V power supply
Electronics Enclosure: IP54.
Input/Output Cable: 5m x 6 core for power in, 4-20mA and pulse output.
Dimensions: 250mm x 48mm x 90mm [10° x 2° x 4°] {electronics + guide rail}.

Figure C-5: Flowmeter Specifications

C.3.2 Temperature Sensors: Platinum Resistance Thermometers (PRT)

Specifications:

- -200°C to $+850^{\circ}\text{C}$ Temperature range
- $\pm 0.1^{\circ}\text{C}$ Accuracy
- 100Ω at 0°C Nominal resistance

C.4 EcoStuxture Setup and Site Navigation

Access through the IP address: <http://18.9.6.83/>

Or desktop workstation (to contact Schneider Electric to obtain the software and license)

1. Log onto MIT VPN via **Global Protect**




Figure C-6: Global Protect VPN and Gateway

2. After connected, choose Gateway as “**GP-BMS-Gateway**”
3. Log into Workstation
 - (a) in username and password; this will be provided by Schneider
 - (b) Type in IP address
 - (c) Hit “LOG ON” and a dialog box will pop up to request your approval, click “Trust Certificate”

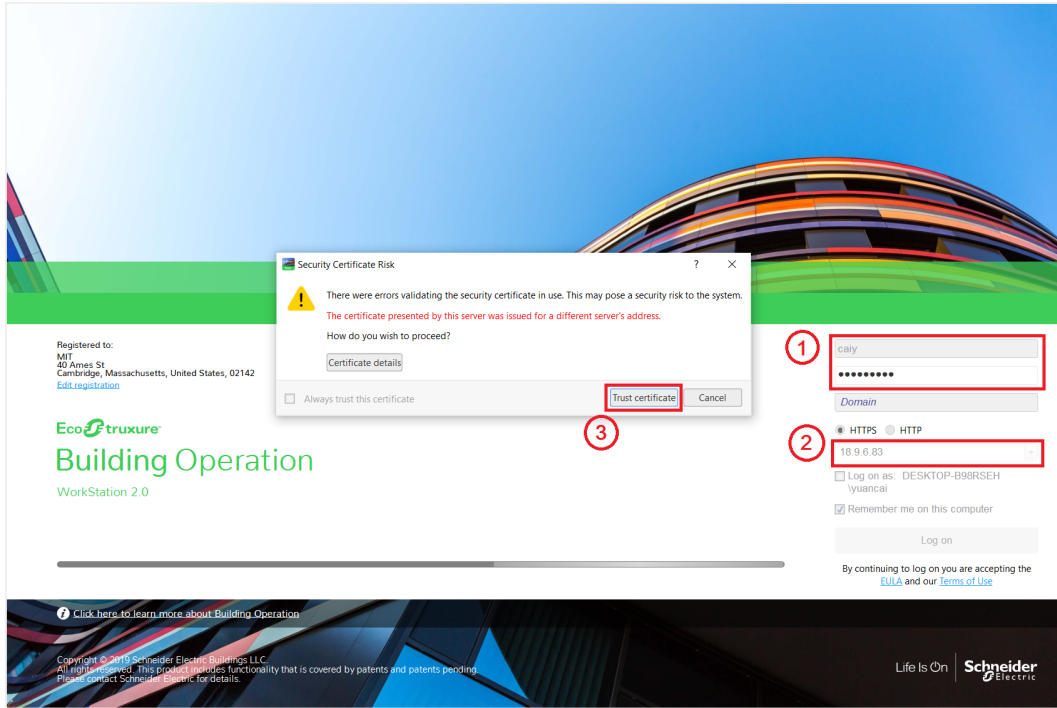


Figure C-7: Workstation login type

4. After login, you will see the main panel (Figure C-8)

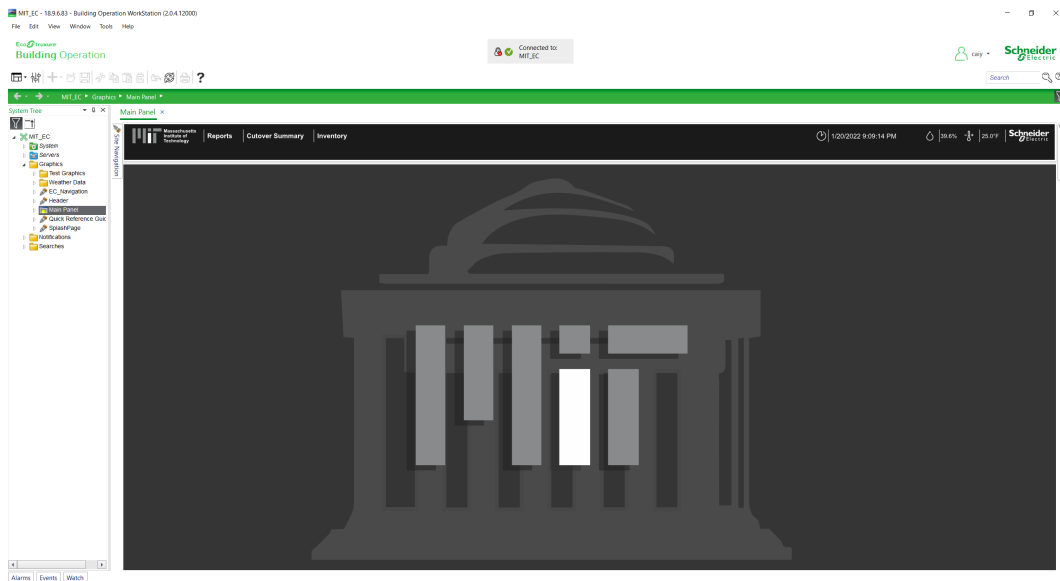


Figure C-8: Main panel

- (a) On the top right corner, click “Site Navigation”, which directs you to all the buildings on campus whose HVAC systems are managed by the EcoStuxture platform (Figure C-9)

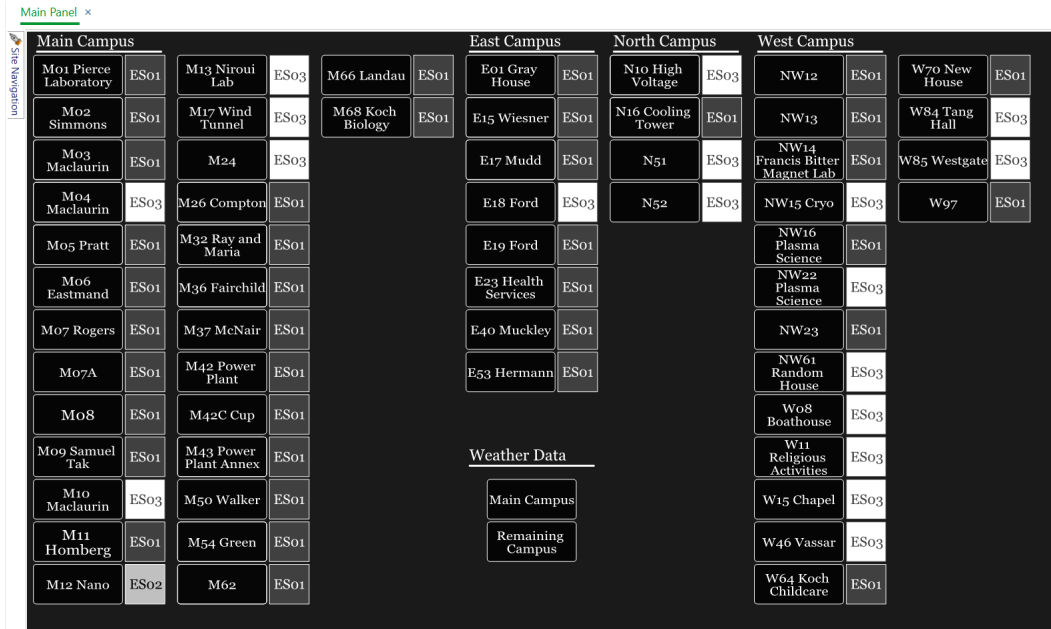


Figure C-9: Site Navigation

- (b) Click the building of interest (i.e. Building 66 [M66]); the left side bar includes an “AI” tab that is the dashboard for setpoint adjustments and heating and cooling loads (chilled beam and FTR) data download; the right side bar allows for an overview of all the floor plans within this building; clicking the floor plan will lead you to the detailed one (Figure C-10)

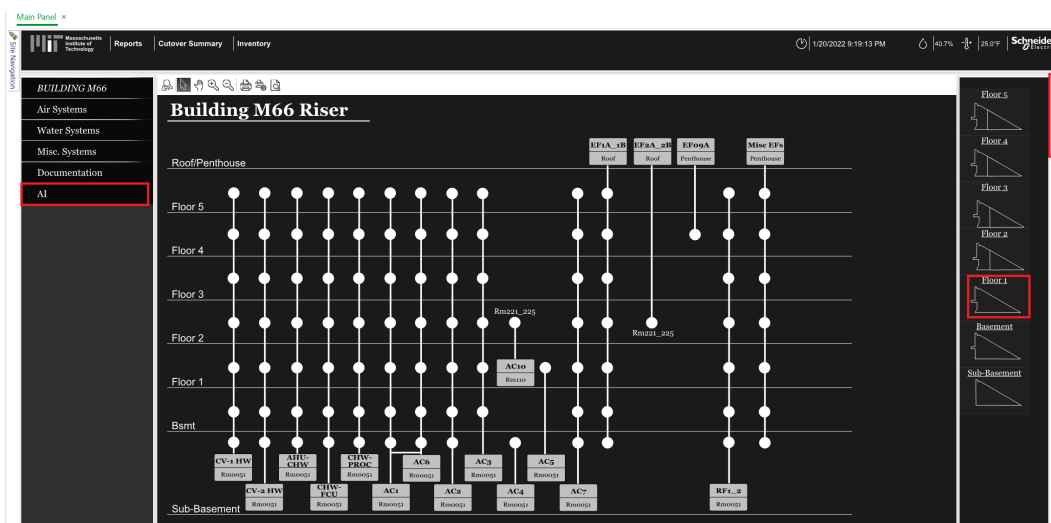


Figure C-10: Floor plan overview, AI project dashboard

- (c) The setpoint adjustment dashboard (Figure C-11) enables you to define setpoints for the next 24 hours at one time; please read the instructions on the left carefully.

Note that the setpoint adjustments are applied to both heating and cooling setpoints; if a varying setpoint deadband is desired, setpoints should be controlled and forced on the room control panel. This will be explained in the following tutorial. The key takeaway is this dashboard will not support a constant cooling setpoint and a changing heating setpoint or vice versa.

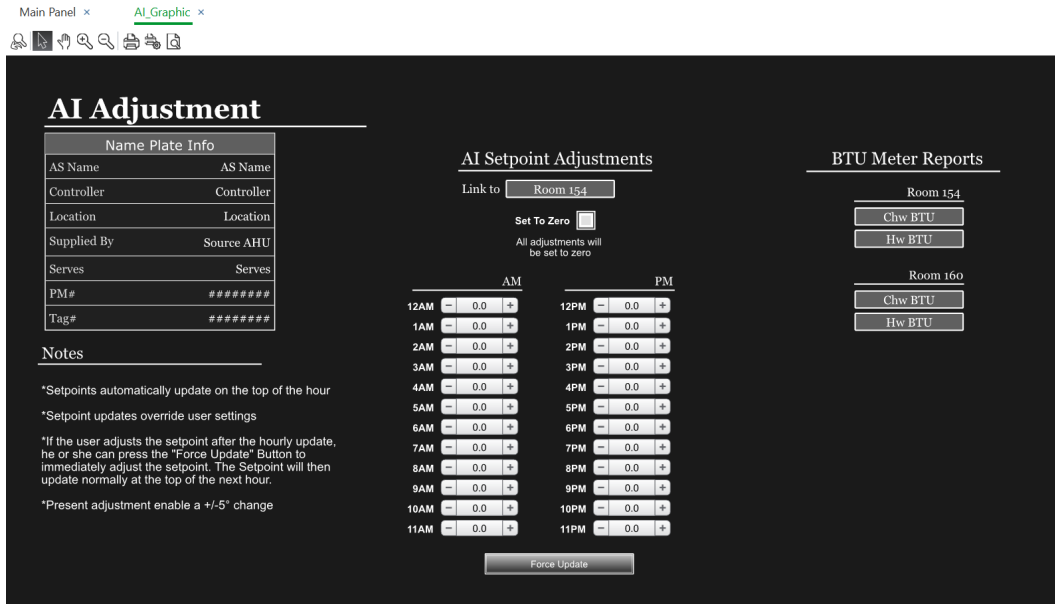


Figure C-11: Setpoint adjustment dashboard and BTU meter reports

- (d) To obtain the BTU meter reports, click the tabs below, which leads you to the new page (Figure C-12). Click “Hide parameters” to unfold the data selection section. Here you can define the start and end report date or specify a date range. Then click “Submit Form” to request the data. ****IMPORTANT****: do not jump to other pages or use any other software while requesting the data, because this will mess up how data are written in the spreadsheet.

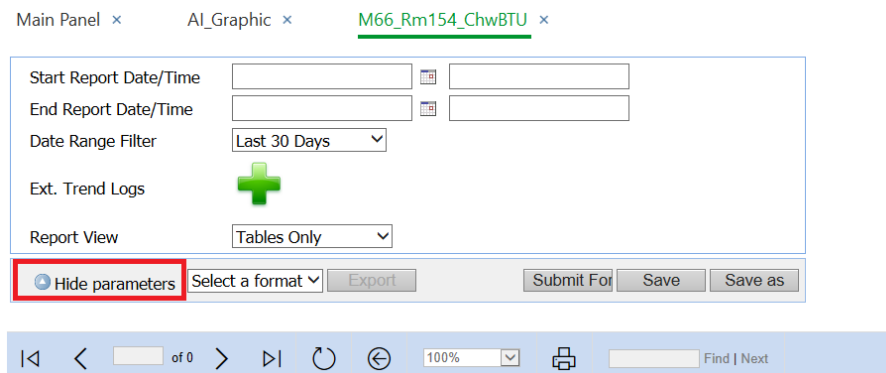


Figure C-12: BTU data download

(e) After data are requested and summary shows up as Figure C-13, “select a format” to export. Repeat the same procedure for other BTU meter data.

Main Panel × AI_Graphic × M66_Rm154_ChwBTU ×

Show parameters Select a format Export Submit For Save Save as

100% Find | Next

Ecostruxure Trend Log Comparison Schneider Electric

Summary	/E/S01/AI Trends/M66 Rm154_ChwBTU_Calculated Power_Tr_AITr (Btu/s)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Flow_Tr_AITr (USgal/min)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Temperature Cold_Tr_AITr (°F)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Temperature Difference_Tr_AITr (°F)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Temperature Hot_Tr_AITr (°F)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Total_Tr_AITr (USgal)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Measured Velocity_Tr_AITr (m/s)	/E/S01/AI Trends/M66 Rm154_ChwBTU_System Type_Tr_AITr (J)	/E/S01/AI Trends/M66 Rm154_ChwBTU_Calculated Energy_Tr_AITr (Btu)
Minimum	0.00	0.00	58.99	0.00	59.58	934671.00	0.00	12.00	577567.69
Maximum	4.11	21.97	74.70	1.43	74.53	941197.00	1.22	12.00	587903.00
Average	0.00	0.19	72.75	0.01	72.62	938494.92	0.01	12.00	584536.70

Report period range: 12/22/2021 12:00:00 AM to 12/12/2022 12:00:00 AM
Report generated: 1/20/2022 9:47:12 PM

Figure C-13: Requested data summary

(f) To enable a constant cooling setpoint and a changing heating setpoint (our preferred control strategy in the heating season), enter the room system control page (Figure C-14) through the floor plan, and force the cooling setpoint by clicking on the temperature. For example, if we want a constant cooling setpoint at 76 °F, force it at 76 °F. Do not force the heating setpoint. Forcing setpoints on the room control pages overrides the setpoint adjustments in the AI page.

M66 Rm154

Name Plate Info

AS Name	M66-AS03
Controller	M66_Rm154
Location	Room 154
Supplied By	####
Serves	Room 154
PM#	N/A
Tag#	####

System Summary

Occupancy	On
Occ Sensor	On
Stand By	Off
Occ Override	Off
Ppl Counter	2.0

Room Temp

Temp	68.5 °F
Rm Temp Base Spt	70.5 °F
Rm Effective Clg Spt	72.0 °F
Rm Effective Htg Spt	69.0 °F
Physical Spt Adj	0.0 Δ°F

Room Humidity

Rh	74.1% Rh
Rm Dewpoint	2.3 °F
Rm CO2	314.4 ppm
CO2 Spt	850.0 ppm

Hw BTU Meter

Calculated Energy	40404996 Btu
Measured Flow	0.0 USgal/min
Measured Total (US Gal)	111826.0 USgal
Measured Temperature Cold	89.0 °F
Measured Temperature Hot	152.4 °F
Measured Temperature Difference	62.8 Δ°F

Chw BTU Meter

Calculated Energy	251158 Btu
Measured Flow	0.0 USgal/min
Measured Total (US Gal)	941197.0 USgal
Measured Temperature Cold	73.4 °F
Measured Temperature Hot	73.6 °F
Measured Temperature Difference	0.0 Δ°F

RmClgSpt Properties

RmClgSpt Value (°F): **Unforce 72**

Room Cooling Setpoint

Reheat Valve Pos: 27.3 %

Cooling Valve Pos: 0.0 %

Chilled Water

Temp	73.4 °F
Critical Spt	55.0 °F
Pump Cmd	Off
Pump Sts	Off

Multiple Chilled Beams in Zone

INFO TREND NOTES

Figure C-14: Forcing setpoints to allow a varying setpoint deadband

C.5 Documentation of BMS and Other Databases Availability for Campus Buildings

Below lists the availability of building management systems and information databases of each campus building. Detailed information can be found in the Git repository: <https://github.com/YuanC233/Thesis-setpoint-optimization.git>.

Main Campus Buildings	KGS	BMS	All Zone	Some Zone	No Zone	PI	No KGS	BMS available but KGS is PI Only
M01	1	1		1		1		
M02	1	1			1	1		
M03	1	1			1	1		
M04	1	1			1	1		
M05	1	1			1	1		
M06	1	1			1	1		
M06C	1	1			1	1		
M07	1	1			1	1		
M07A							1	
M08	1	1			1	1		
M09	1	1			1	1		
M10	1	1			1	1		
M11	1	1			1			
M12	1					1		
M13	1	1	1			1		
M14	1					1		1
M16	1	1			1	1		
M18	1	1	1			1		
M24	1	1			1	1		
M26	1	1			1	1		
M31	1	1			1	1		
M32	1	1		1		1		
M33	1	1			1	1		
M34	1					1		1
M35	1	1			1	1		
M36	1	1			1	1		
M37	1	1			1	1		
M38	1	1			1	1		
M39	1					1		1
M42	1	1			1			
M43							1	
M46	1	1	1			1		
M48	1	1			1	1		
M50	1					1		1
M54	1					1		1
M56	1	1			1	1		

Figure C-15: BMS and other databases availability for campus buildings

Main Campus Buildings	KGS	BMS	All Zone	Some Zone	No Zone	PI	No KGS	BMS available but KGS is PI Only
M57							1	
M62	1					1		1
M64	1					1		
M66	1	1	1			1		
M68	1	1		1		1		
M76	1	1	1			1		
 								
E01	1					1		1
E02	1	1			1	1		
E14	1	1	1			1		
E15	1					1		1
E17	1	1			1	1		
E18	1	1			1	1		
E19	1	1			1	1		
E23	1	1			1	1		
E25	1	1		1		1		
E40	1					1		1
E51	1	1			1	1		
E52	1	1	1			1		
E53	1					1		1
E55	1	1			1	1		
E60	1	1	1			1		
E62	1	1	1			1		
 								
W01	1	1		1		1		
W02	1					1		
W04							1	
W07	1	1			1			
W08							1	
W11							1	
W13	1					1		
W15	1					1		1
W16	1	1	1			1		
W20	1					1		1
W31	1	1			1			
W32							1	

Figure C-16: BMS and other databases availability for campus buildings (continued)

<u>Main Campus Buildings</u>	<u>KGS</u>	<u>BMS</u>	<u>All Zone</u>	<u>Some Zone</u>	<u>No Zone</u>	<u>PI</u>	<u>No KGS</u>	<u>BMS available but KGS is PI Only</u>
W34	1					1		1
W35	1	1	1			1		
W51							1	
W53a	1	1			1			
W59	1	1			1			
W61	1					1		1
W64							1	
W70	1	1			1	1		
W71	1	1			1			
W79	1	1	1			1		
W84	1	1			1			
W85	1					1		1
W91	1	1			1			
W92							1	
W98	1	1			1			
N10							1	
N16	1	1			1			
N51	1	1			1			
N52	1	1			1			
NW10							1	
NW12							1	
NW13							1	
NW14							1	
NW14T	1	1			1	1		
NW15							1	
NW16							1	
NW17							1	
NW21	1					1		1
NW22	1	1			1			
NW23	1	1	1			1		
NW30	1	1			1	1		
NW35	1	1		1		1		
NW61							1	
NW86	1	1			1			

Figure C-17: BMS and other databases availability for campus buildings (continued)

<u>Main Campus Buildings</u>	<u>KGS</u>	<u>BMS</u>	<u>All Zone</u>	<u>Some Zone</u>	<u>No Zone</u>	<u>PI</u>	<u>No KGS</u>	<u>BMS available but KGS is PI Only</u>
Totals>	85	65	13	6	46	70	19	16

Figure C-18: BMS and other databases availability for campus buildings (continued)

Appendix D

Source Code

The Python and MATLAB source code for the simulations performed in Chapters 3-5 are included in GitHub Repo: <https://github.com/YuanC233/Thesis-setpoint-optimization.git>

Appendix E

Documentation of EnergyPlus Simulation

The simulation IDF files are included in the GitHub Repo: <https://github.com/YuanC23/3/Thesis-setpoint-optimization.git>

E.1 Room 154 Model Parameters




Figure E-1: [Room 154] EnergyPlus Simulation model

Room dimension: $24' \times 32' \times 12' = 7.4\text{m} \times 9.7\text{m} \times 3.66\text{m}$

Walls

Exterior: south wall only

Interior: east, west and north walls, roof and floor (has basement) (adiabatic)

Material: M11 300mm concrete (in reality 1m thick)

Window

Dimension: 3m × 8m

Material: single pane (low e, tinted)

Window faces 20 degrees east of due south

Schedules

Lighting schedule: max 700W

Occupancy schedule: 0 (no occupancy due to COVID)

Setpoint schedule:

- Cooling season: 4/1 – 10/15; cooling setpoint: 21.4 °C; heating setpoint: 10 °C
- Heating season: 1/1 – 3/31, 10/16 – 12/31; heating setpoint: 21.4 °C; cooling setpoint: 40 °C

HVAC

Packaged DX cooling with gas heat

Design day (.ddy file)

- Heating season: 99%
- Cooling: 1%

Others

Shading control with blinds

24' × 32' × 12' = 7.4m × 9.7m × 3.66m

Bibliography

- [1] Buildings - alliance to save energy. <https://www.ase.org/initiatives/buildings>. Online available.
- [2] Dynamic Pricing and Smart Grid. www.elevateenergy.org. Online available.
- [3] Energy efficiency: Buildings. <https://www.iea.org/topics/energyefficiency/buildings/>. Online available.
- [4] FAQs: Locational Marginal Pricing. <https://www.iso-ne.com/participate/support/faq/lmp>. Online available.
- [5] Frequently Asked Questions (FAQs). <https://www.eia.gov/tools/faqs/faq.php?id=105&t=3>. Online available.
- [6] Hourly electricity consumption varies throughout the day and across seasons. <https://www.eia.gov/todayinenergy/detail.php?id=42915#>. Online available.
- [7] Ontario Energy Board. www.oeb.ca. Online available.
- [8] Residential Energy Incentive Programs. https://www.pge.com/en_US/residential/save-energy-money/savings-solutions-and-rebates/demand-response/demand-response.page. Online available.
- [9] Thermostat. <https://facilities.ucdavis.edu/energy-engineering/thermostat>. Online available.
- [10] WattTime. <https://www.watttime.org/>. Online available.
- [11] Ratnadip Adhikari and Ramesh K Agrawal. An introductory study on time series modeling and forecasting. *arXiv preprint arXiv:1302.6613*, 2013.
- [12] U.S. Energy Information Administration. Annual Energy Outlook 2021. Technical report, U.S. Energy Information Administration, 2021.
- [13] Abdul Afram and Farrokh Janabi-Sharifi. Theory and applications of hvac control systems—a review of model predictive control (mpc). *Building and Environment*, 72:343–355, 2014.
- [14] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.

- [15] Jonghoon Ahn and Soolyeon Cho. Dead-band vs. machine-learning control systems: Analysis of control benefits and energy efficiency. *Journal of Building Engineering*, 12:17–25, 2017.
- [16] Khalid Al-Jabery, Don C Wunsch, Jinjun Xiong, and Yiyu Shi. A novel grid load management technique using electric water heaters and q-learning. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 776–781. IEEE, 2014.
- [17] Emilio Ancillotti, Raffaele Bruno, and Marco Conti. Smoothing peak demands through aggregate control of background electrical loads. In *ISGT 2014*, pages 1–5. IEEE, 2014.
- [18] ANSI Ashrae and ASHRAE Standard. 62.1. 2019, ventilation for acceptable indoor air quality. *American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA*, 2019.
- [19] M Asif, T Muneer, and R Kelley. Life cycle assessment: A case study of a dwelling home in scotland. *Building and environment*, 42(3):1391–1394, 2007.
- [20] Amotz Bar-Noy, Matthew P Johnson, and Ou Liu. Peak shaving through resource buffering. In *International workshop on Approximation and Online Algorithms*, pages 147–159. Springer, 2008.
- [21] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht. Smartcap: Flattening peak electricity demand in smart homes. In *2012 IEEE International Conference on Pervasive Computing and Communications*, pages 67–75. IEEE, 2012.
- [22] Enda Barrett and Stephen Linder. Autonomous hvac control, a reinforcement learning approach. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 3–19. Springer, 2015.
- [23] Stuart Batterman. Review and extension of co2-based methods to determine ventilation rates with application to school classrooms. *International journal of environmental research and public health*, 14(2):145, 2017.
- [24] David Henry Blum. *Improving the use of commercial building HVAC systems for electric grid ancillary services*. PhD thesis, Massachusetts Institute of Technology, Department of Architecture, 2016.
- [25] David P Chassin, Jakob Stoustrup, Panajotis Agathoklis, and Nedjib Djilali. A new thermostat for real-time price demand response: Cost, comfort and energy impacts of discrete-time control without deadband. *Applied Energy*, 155:816–825, 2015.
- [26] Kein Huat Chua, Yun Seng Lim, and Stella Morris. Battery energy storage system for peak shaving and voltage unbalance mitigation. *International Journal of Smart Grid and Clean Energy*, 2(3):357–363, 2013.
- [27] JA Crabb, N Murdoch, and JM Penman. A simplified thermal response model. *Building Services Engineering Research and Technology*, 8(1):13–19, 1987.

- [28] Drury B Crawley, Linda K Lawrie, Frederick C Winkelmann, Walter F Buhl, Y Joe Huang, Curtis O Pedersen, Richard K Strand, Richard J Liesen, Daniel E Fisher, Michael J Witte, et al. Energyplus: creating a new-generation building energy simulation program. *Energy and buildings*, 33(4):319–331, 2001.
- [29] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 2016. To appear.
- [30] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [31] Richard E Edwards, Joshua New, Lynne E Parker, Borui Cui, and Jin Dong. Constructing large scale surrogate models from big data and artificial intelligence. *Applied energy*, 202:685–699, 2017.
- [32] Peter G Ellis, Paul A Torcellini, and D Crawley. Simulation of energy management systems in energyplus. 2008.
- [33] SL Englander and LK Norford. Variable speed drives: improving energy consumption modeling and savings analysis techniques. *Proc. A CEEE 1992 Summer Study on Energy Efficiency in Buildings*, 1992.
- [34] Varick L Erickson and Alberto E Cerpa. Occupancy based demand response hvac control strategy. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 7–12, 2010.
- [35] Torgeir Ericson. Direct load control of residential water heaters. *Energy Policy*, 37(9):3502–3512, 2009.
- [36] Mohammad Esrafilian-Najafabadi and Fariborz Haghighat. Occupancy-based hvac control systems in buildings: A state-of-the-art review. *Building and Environment*, page 107810, 2021.
- [37] Stefan Feuerriegel and Dirk Neumann. Measuring the financial impact of demand response for electricity retailers. *Energy Policy*, 65:359–368, 2014.
- [38] Nelson Fumo, Pedro Mago, and Rogelio Luck. Methodology to estimate building energy consumption using energyplus benchmark models. *Energy and Buildings*, 42(12):2331–2337, 2010.
- [39] Everette S Gardner Jr. Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28, 1985.
- [40] Pedro A Gonzalez and Jesus M Zamarreno. Prediction of hourly energy consumption in buildings based on a feedback artificial neural network. *Energy and buildings*, 37(6):595–601, 2005.
- [41] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.

- [42] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [43] He Hao, Borhan M Sanandaji, Kameshwar Poolla, and Tyrone L Vincent. Aggregate flexibility of thermostatically controlled loads. *IEEE Transactions on Power Systems*, 30(1):189–198, 2014.
- [44] Andrew C Harvey and Simon Peters. Estimation procedures for structural time series models. *Journal of forecasting*, 9(2):89–108, 1990.
- [45] Trevor Hastie and Robert Tibshirani. Generalized additive models: some applications. *Journal of the American Statistical Association*, 82(398):371–386, 1987.
- [46] IEA. Technology Roadmap - Smart Grids. Technical report, IEA, 2011.
- [47] Kevin J Kircher, Yuan Cai, Leslie K Norford, and Steven B Leeb. Controlling big, diverse, nonlinear load aggregations for grid services by adjusting device setpoints. In *IEEE Conference on Decision and Control*. IEEE, 2021.
- [48] Rick Kramer, Jos Van Schijndel, and Henk Schellen. Simplified thermal and hygric building models: A literature review. *Frontiers of architectural research*, 1(4):318–325, 2012.
- [49] Matt Kraning, Eric Chu, Javad Lavaei, Stephen P Boyd, et al. *Dynamic network energy management via proximal message passing*. Citeseer, 2014.
- [50] SC Lee, SJ Kim, and SH Kim. Demand side management with air conditioner loads based on the queuing system model. *IEEE Transactions on Power Systems*, 26(2):661–668, 2010.
- [51] Yashen Lin, Timothy Middelkoop, and Prabir Barooah. Issues in identification of control-oriented thermal models of zones in multi-zone buildings. In *2012 IEEE 51st IEEE conference on decision and control (CDC)*, pages 6932–6937. IEEE, 2012.
- [52] Ning Lu and Srinivas Katipamula. Control strategies of thermostatically controlled appliances in a competitive electricity market. In *IEEE Power Engineering Society General Meeting, 2005*, pages 202–207. IEEE, 2005.
- [53] Jingran Ma, S Joe Qin, Bo Li, and Tim Salsbury. *Economic model predictive control for building energy systems*. IEEE, 2011.
- [54] Yudong Ma, Anthony Kelman, Allan Daly, and Francesco Borrelli. Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments. *IEEE control systems magazine*, 32(1):44–64, 2012.
- [55] Mehdi Maasoumy and Alberto Sangiovanni-Vincentelli. Total and peak energy consumption minimization of building hvac systems using model predictive control. *IEEE Design & Test of Computers*, 29(4):26–35, 2012.
- [56] Karel Mařík, Jiří Rojíček, Petr Stluka, and Jiří Vass. Advanced hvac control: Theory vs. reality. *IFAC Proceedings Volumes*, 44(1):3108–3113, 2011.

- [57] Andrei Marinescu, Ivana Dusparic, Adam Taylor, Vinny Cahill, and Siobhán Clarke. P-marl: Prediction-based multi-agent reinforcement learning for non-stationary environments. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1897–1898, 2015.
- [58] Clayton Miller, Pandarasamy Arjunan, Anjukan Kathirgamanathan, Chun Fu, Jonathan Roth, June Young Park, Chris Balbach, Krishnan Gowri, Zoltan Nagy, Anthony D Fontanini, et al. The ashrae great energy predictor iii competition: Overview and results. *Science and Technology for the Built Environment*, 26(10):1427–1447, 2020.
- [59] Truong X Nghiem, Madhur Behl, Rahul Mangharam, and George J Pappas. Green scheduling of control systems for peak demand reduction. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 5131–5136. IEEE, 2011.
- [60] HyungSeon Oh and Robert J Thomas. Demand-side bidding agents: Modeling and simulation. *IEEE Transactions on Power Systems*, 23(3):1050–1056, 2008.
- [61] Frauke Oldewurtel, David Sturzenegger, and Manfred Morari. Importance of occupancy information for building climate control. *Applied energy*, 101:521–532, 2013.
- [62] Frauke Oldewurtel, Andreas Ulbig, Alessandra Parisio, Göran Andersson, and Manfred Morari. Reducing peak electricity demand in building climate control using real-time pricing and model predictive control. In *49th IEEE conference on decision and control (CDC)*, pages 1927–1932. IEEE, 2010.
- [63] M Pak, Adrian Smith, and Gordon Gill. Ladybug : a parametric environmental plugin for grasshopper to help designers create an environmentally-conscious design. 2013.
- [64] Peter Palensky and Dietmar Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE transactions on industrial informatics*, 7(3):381–388, 2011.
- [65] Cristian Perfumo, Julio H Braslavsky, and John K Ward. Model-based estimation of energy savings in load control events for thermostatically controlled loads. *IEEE Transactions on Smart Grid*, 5(3):1410–1420, 2014.
- [66] Santosh Philip. Energyplus eppy, version 0.5.57. <https://pypi.org/project/eppy/>.
- [67] Kyu-Nam Rhee, Sun-Ho Choi, and Agisilaos Tsouvalas. Experimental investigation on the control performance of an active chilled beam system under dynamic cooling loads. *Applied Thermal Engineering*, 194:117069, 2021.
- [68] Edwin Rodriguez-Ubinas, Claudio Montero, María Porteros, Sergio Vega, Iñaki Navarro, Manuel Castillo-Cagigal, Eduardo Matallanas, and Alvaro Gutiérrez. Passive design strategies and performance of net energy plus houses. *Energy and buildings*, 83:10–22, 2014.
- [69] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019.

- [70] Kurt Roth, John Dieckmann, Robert Zogg, and James Brodrick. Chilled beam cooling. *Ashrae journal*, 49(9):84, 2007.
- [71] Nerea Ruiz, Iñigo Cobelo, and José Oyarzabal. A direct load control model for virtual power plant management. *IEEE Transactions on Power Systems*, 24(2):959–966, 2009.
- [72] Alaa Sagheer and Mostafa Kotb. Time series forecasting of petroleum production using deep lstm recurrent networks. *Neurocomputing*, 323:203–213, 2019.
- [73] Borhan M Sanandaji, He Hao, and Kameshwar Poolla. Fast regulation service provision via aggregation of thermostatically controlled loads. In *2014 47th Hawaii International Conference on System Sciences*, pages 2388–2397. IEEE, 2014.
- [74] Borhan M Sanandaji, Tyrone L Vincent, and Kameshwar Poolla. Ramping rate flexibility of residential hvac loads. *IEEE Transactions on Sustainable Energy*, 7(2):865–874, 2015.
- [75] Spencer C Shabshab, Peter A Lindahl, Steven B Leeb, and J Kendall Nowocin. Autonomous demand smoothing for efficiency improvements on military forward operating bases. *IEEE Transactions on Power Delivery*, 35(5):2243–2251, 2020.
- [76] Elham Shirazi and Shahram Jadid. Cost reduction and peak shaving through domestic load shifting and ders. *Energy*, 124:146–159, 2017.
- [77] Pierluigi Siano. Demand response and smart grids—a survey. *Renewable and sustainable energy reviews*, 30:461–478, 2014.
- [78] Servet Soyguder, Mehmet Karakose, and Hasan Alli. Design and simulation of self-tuning pid-type fuzzy adaptive control for an expert hvac system. *Expert systems with applications*, 36(3):4566–4573, 2009.
- [79] Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.
- [80] Tianshu Wei, Yanzhi Wang, and Qi Zhu. Deep reinforcement learning for building hvac control. In *Proceedings of the 54th annual design automation conference 2017*, pages 1–6, 2017.
- [81] White-House. Fact sheet: President Biden sets 2030 greenhouse gas pollution reduction target aimed at creating good-paying union jobs and securing U.S. leadership on clean energy technologies, 2021.
- [82] Liu Yang, Haiyan Yan, and Joseph C Lam. Thermal comfort and building energy consumption implications—a review. *Applied Energy*, 115:164–173, 2014.
- [83] Jinghua Yu, Changzhi Yang, Liwei Tian, and Dan Liao. A study on optimum insulation thicknesses of external walls in hot summer and cold winter zone of China. *Applied Energy*, 86(11):2520–2529, 2009.
- [84] Liang Yu, Yi Sun, Zhanbo Xu, Chao Shen, Dong Yue, Tao Jiang, and Xiaohong Guan. Multi-agent deep reinforcement learning for hvac control in commercial buildings. *IEEE Transactions on Smart Grid*, 12(1):407–419, 2020.

- [85] J Zarnikau. Chapter 8: Demand participation in restructured markets, competitive electricity markets: Design, implementation, performance, 2008.
- [86] Z. Zhang, K.J. Kircher, Y. Cai, J.G. Brearly, D. Birge, and L.K. Norford. Mitigating peak load and heat stress under heat waves by optimizing thermostat setpoint and fan speed schedules. *Science and Technology for the Built Environment*, 2021. Submitted.
- [87] Jie Zhao, Khee Poh Lam, and B Erik Ydstie. Energyplus model-based predictive control (epmpc) by using matlab/simulink and mle+. 2013.
- [88] Jie Zhao, Khee Poh Lam, B Erik Ydstie, and Vivian Loftness. Occupant-oriented mixed-mode energyplus predictive control simulation. *Energy and Buildings*, 117:362–371, 2016.