Economic Optimization of Large-Scale Commercial Building Heating, Ventilation, and Air Conditioning Systems

by

Nishith R. Patel

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Chemical Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2018

Date of final oral examination: 08/10/2018

The dissertation is approved by the following members of the Final Oral Committee:

James B. Rawlings, Professor, Chemical and Biological Engineering
Christos T. Maravelias, Professor, Chemical and Biological Engineering
Victor M. Zavala, Assistant Professor, Chemical and Biological Engineering
Ross E. Swaney, Associate Professor, Chemical and Biological Engineering
Jeffrey T. Linderoth, Professor, Industrial and Systems Engineering

To my family.

ACKNOWLEDGMENTS

This dissertation would not be possible without the help of many people.

First, I would like to thank my advisor, Professor James B. Rawlings, for his endless guidance on research and life. I would always walk away from a meeting with Prof. Rawlings with a new perspective on how to approach the problem at hand. He has an uncanny ability to extract out the essence of a complex problem and boil it down to the fundamentals. Even outside of research, he has made the graduate school years enjoyable by exposing me to cross-country skiing in Madison for which I am grateful.

I thank Professors Christos T. Maravelias and Victor M. Zavala for their collaborations on my research project. I learned a lot on the subjects of mixed-integer optimization and stochastic programming from my interactions with them during my time at UW-Madison. I would also like to thank Professors Ross E. Swaney and Jeffrey T. Linderoth for taking the time to serve on my defense committee and to provide valuable insight.

I would like to acknowledge Bob Turney from Johnson Controls, Inc. for bringing the research problem statement presented in this work to our group. I also thank the fellow JCI employees with whom I have had fruitful discussions related to commercial building systems: Kirk Drees, Mike Wenzel, and Matt Ellis.

Over the years, I have had the pleasure of working with talented colleagues in the Rawlings group. I thank former members Megan, Ankur, and Luo for helping to get me up and running with the group's software and hardware tools. I thank Michael, Min, Travis, and Doug for being there alongside me during most of my journey in graduate school. I appreciated your knowledge and help during the serious times as well as the moments of laughter during the lighter times. I thank Mary Heimbecker as well for her administrative help and sweet treats. I wish new members Koty and

Pratyush all the best as they begin their journey at UC Santa Barbara. I also thank members of the Maravelias and Zavala research groups for helpful conversations along the way.

For software, I used MATLAB and GNU Octave for performing the HVAC system calculations presented in this dissertation. Most of the optimization problems were solved using Gurobi, and those in Chapter 7 were solved using CVX. This document was prepared using LaTeX. I thank Will Benton whose publicly available UW dissertation template on GitHub I used as a starting point. For plotting results, I used the Python package Matplotlib. I used the LaTeX package TikZ to produce the drawings in this work. I am grateful to the developers for these software tools for making my life easier. I also want to thank Michael Risbeck in our group whose skeleton repository structure saved me valuable time during the document production phase.

Finally, but certainly not least, I want to thank my family for their unconditional support through the good and bad times and also for letting me venture this far north by myself. See, I told you guys I packed enough jackets to make it through the five ice ages with almost all of my extremities intact.

Nishith Patel
Madison, WI
August 2018

CONTENTS

LIST OF TABLES

LIST OF FIGURES

NOMENCLATURE

Notation

n_b	number of buildings
n_{bz}	number of zones in the building
$u \in \mathbb{R}^m$	vector of inputs
$x \in \mathbb{R}^n$	vector of states
$y \in \mathbb{R}^p$	vector of measurements
$d \in \mathbb{R}^{n_d}$	vector of disturbances

Relevant Sets and Indices

$b \in \mathbb{B}$	buildings (high-level)
$i \in \mathbb{I}_b$	zones in the building (low-level airside)
$k \in \mathbb{K}$	discrete time

Variables

T	temperature
T_a	ambient temperature
$T_{sp,i}$	zone temperature setpoint
$\dot{Q}_{other}$	external load, radiation, or disturbance
$\dot{Q}_c$	cooling rate delivered
Q_{total}	total amount of cooling delivered
$\dot{Q}_{HVAC}$	cooling rate from HVAC system
$\dot{Q}_{peak}$	peak HVAC system cooling rate
$\dot{W}_{peak}$	peak power usage
s	amount of cooling potential in storage tank
$\dot{Q}_{storage}$	cooling rate delivered to storage tank
$\dot{Q}_{HighLevel}$	cooling rate from high-level problem
$T_{HighLevel}$	building temperature from high-level problem
ε_i	tracking error

$\bar{\varepsilon}_i$	integral of tracking error
$\nu_{air,i}$	air flow rate in air handler unit

Parameters

Δ	sample time of controller
N	horizon length
c_k	cost of electricity at time k
$\tilde{c}_k$	cost of cooling at time k
c_{peak}	peak demand charge
C	thermal capacitance
H	scaled ambient heat transfer coefficient
β_{ij}	scaled heat transfer coefficient between zones
$K_{c,i}$	scaled zone PI controller gain
$\tau_{I,i}$	integral time constant for zone PI controller
$\dot{Q}_{ss,i}$	steady-state rate of cooling
η_{air}	inverse COP of air handler units
η_{batt}	round-trip battery efficiency
η_{HVAC}	inverse COP of HVAC (waterside) system
$\eta_{storage}$	efficiency of storage tank
η_{tot}	inverse COP of aggregate airside/waterside system
$\dot{Q}_{HVAC,max}$	max cooling capacity of HVAC system
$\dot{Q}_{peak,past}$	peak cooling rate previously achieved
$\dot{W}_{peak,past}$	peak power usage previously achieved
σ	decay constant for storage tank
s_{max}	max cooling capacity of storage tank
T_{min}	lower bound of comfort region
T_{max}	upper bound of comfort region
$\dot{Q}_{c,i,min}$	min cooling rate
$\dot{Q}_{c,i,max}$	max cooling rate
μ	penalty on energy usage

α	rate of change penalty
β	comfort zone constraint violation penalty
γ	high-level heat duty constraint violation penalty
μ	total amount of cooling delivered penalty
$p_{i,j}$	rate of change of the input from time j to time i
p_k	comfort zone constraint violation at time k
$p_{Q,k}$	high-level heat duty constraint violation at time k
$\nu_{air,ss,i}$	steady-state air flow rate in air handler unit
$\tilde{K}_{c,i}$	PI controller gain
U_i	heat transfer coefficient in air handler unit
A_i	area of heat transfer in air handler unit
HX_i	air handler unit
T_{CW}	chilled water temperature

ACRONYMS

AHU	air-handler unit
BAS	building automation system
COP	coefficient of performance
EMPC	economic model predictive control
FR	frequency regulation
HRC	heat recovery chiller
HVAC	heating, ventilation, and air conditioning
LP	linear program
MILP	mixed-integer linear program
MIP	mixed-integer program
MPC	model predictive control
NLP	nonlinear program
OCP	optimal control problem
PI	proportional-integral
PID	proportional-integral-derivative
PRBS	pseudo-random, binary signal
QP	quadratic program
RBC	rule-based control
RTU	roof-top unit
SESI	Stanford Energy System Innovations
SI	system identification
TES	thermal energy storage
VRF	variable refrigerant flow

ABSTRACT

Heating, ventilation, and air conditioning (HVAC) systems of commercial buildings have received significant attention recently due to the large amounts of energy consumed in these applications. Substantial energy cost savings are possible for HVAC systems in utility markets with time-varying price structures through load shifting by using energy storage resources. While model predictive control has been recommended as a strategy to attain these savings, wide-scale implementation has not yet occurred. In large-scale commercial applications with hundreds of zones, solving a single comprehensive mixed-integer optimization problem online for MPC in real time is not feasible. In this dissertation, a hierarchical decomposition of the overall problem is proposed to address this gap. The proposed hierarchical decomposition uses aggregate models of the airside and waterside problems to reduce complexity, and iterations are not required between the subsystem controllers. Energy cost savings between 5% and 40% can be achieved with this control system based on the amount of energy storage and the economic incentives for load shifting.

The decomposition framework is extended to handled systems with embedded battery units. Working towards the long-term objective of implementation, methods for model development are also discussed. Models and data are developed for a large-scale case study loosely based on the Stanford University central HVAC plant. The parameters from the case study are made publicly available for researchers in the HVAC community who would like to design other control architectures and evaluate their performance. The case study is also used to demonstrate that the proposed decomposition can be solved in real time. Finally, a procedure for tracking zone regions using MPC is outlined. Nominal stability for this controller is proven, and its benefits in control performance over heuristic strategies is exemplified via simulation studies.

1 INTRODUCTION

In 2017, residential and commercial buildings were responsible for almost 40% of the total energy consumption in the United States [23]. Commercial buildings alone account for roughly $200 billion a year in primary energy expenditures. Moreover, the Energy Information Administration projects that these primary energy consumption values will increase by 22% between 2009 and 2035 [15]. Average energy prices, on the other hand, have been and are expected to remain relatively stable; therefore, the amount spent on energy in commercial buildings, which is already significant, is going to continue increasing for the foreseeable future.

A large portion of these energy expenditures are due to space heating and cooling. As a result, heating, ventilation, and air conditioning (HVAC) systems have become a major focus of efforts aimed at the development of new energy cost minimization strategies. The presence of time-of-usage pricing structures provides an economic incentive to HVAC customers to earn cost savings through optimization of their equipment and resources. Despite researchers suggesting model predictive control (MPC) as a good candidate for HVAC control to achieve these energy cost savings, better fundamental understanding and implementation strategies are required before this technology can be successfully deployed in practice. The primary research objective of this dissertation is to work towards the design of an industrially implementable, economically optimal HVAC control system.

In this chapter, the economic incentives for HVAC applications are described. The capabilities of MPC to realize these savings are stated. Some barriers to the widespread deployment of control technology are listed. Finally, the research objectives are outlined.

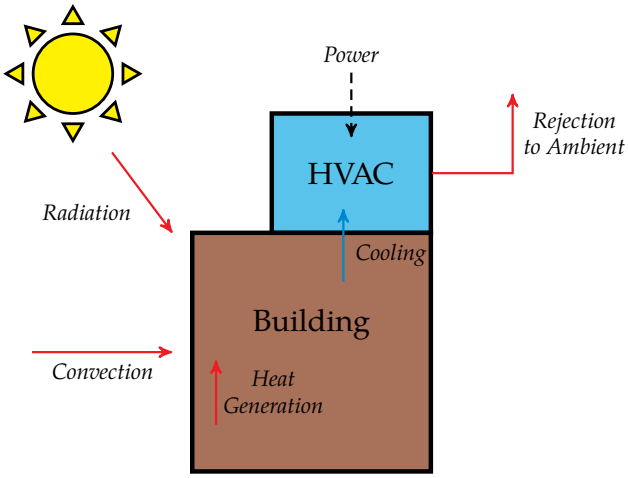

Figure 1.1: Flow of heat in buildings

1.1 Load Shifting

In the summer, several loads are placed on buildings, including radiation from the sun, convection with the ambient, and internal heat generation due to lighting and occupancy. The flow of heat in buildings is depicted in Figure 1.1. The role of the HVAC system is to reject that heat back to the ambient using power purchased from the electricity market.

Currently, most commercial buildings rely on temperature controllers whose goal is to converge to the desired temperature setpoint and stay there within some tolerance. In fact, HVAC systems use simple on/off and proportional-integral-derivative (PID) controllers for this purpose [2, 31]. These controllers are widely used due to their simplicity and robustness [87]. The setpoints given to these thermostats serve as a measure of comfort for occupants in the building. A better goal is to consider economics and minimize total energy cost.

One key opportunity for decreasing energy costs lies in the presence of time-varying prices in utility markets. This type of pricing structure is used by power companies due to the nature of the load placed on them by consumers. Figure 1.2

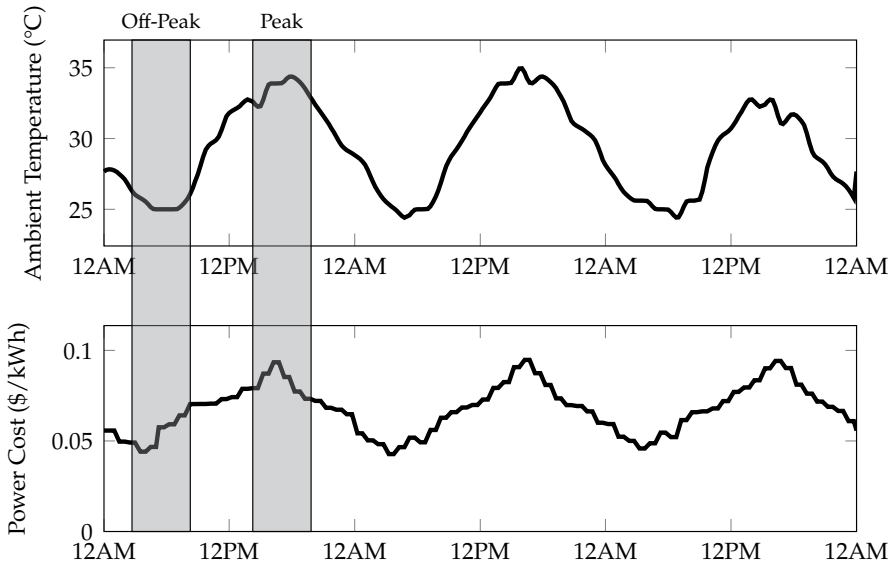

Figure 1.2: Peak and off-peak periods for electricity pricing

depicts typical weather and pricing data for a three-day period in the summer. The weather data was obtained for a city in the southern United States; the temperature is highest in the afternoon hours and lowest in the early morning hours, which corresponds well to the actual heat load placed on HVAC systems. The electricity pricing data was obtained from Johnson Controls, Inc. Note the time-of-usage prices follow a similar pattern: prices are higher in the afternoon hours, referred to as *peak periods*, when the load is greatest on the power companies and prices are lower in the early morning hours, during the *off-peak periods*. Power companies enforce this structure to incentivize more electricity usage during the off-peak periods and less during the peak periods since they operate most efficiently at constant load. Power companies are sized according to their maximum load, hence most of the time, they are running less efficiently below capacity.

From an economic standpoint for the consumers, it is cheaper to run the HVAC systems as much as possible during the night when electricity prices are lower than

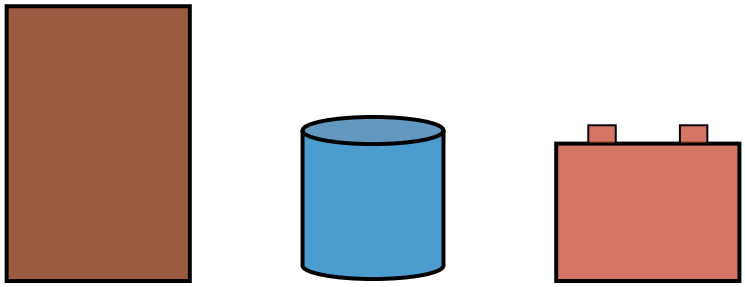

Figure 1.3: Various methods of energy storage

during peak hours. Less power is then required during the afternoon peak hours when demand surges and costs increase. This phenomenom is known as *load shifting*. This movement of the power load reduces the load on power plants during peak hours, so they can operate more efficiently. If power plants waste less energy by operating at higher efficiency, then their total primary energy consumption decreases significantly. In order to move this load, a method of storing the energy is required.

1.2 Methods of Energy Storage

Various methods of energy storage can be used for load shifting, as shown in Figure 1.3. Two commonly used methods include thermal energy storage (TES) and electrical storage.

A TES system can be categorized as "active" or "passive" [29]. In active TES, energy is concentrated in a solid or liquid medium that can be stored at high efficiency (e.g., chilled water or ice tanks). While in passive TES, the mass of the building is used as a natural source of thermal capacitance. In general, a combination of both active and passive TES is the most effective control strategy [29, 32, 33, 41]. Cooling a building or charging active TES tank at night has the added benefit of allowing the chillers to operate more efficiently, since the cooling water temperature is lower, and

thus, the coefficient of performance (COP) is higher. Similar considerations can be applied to heating buildings during the winter [30], however the focus of this work is cooling during the summer due to the variable pricing in electricity markets.

There are also several options for electrical storage with recent advancements in battery technology [99]. With electrical storage, the power purchased during off-peak periods is stored directly in batteries which are discharged when power is needed by the HVAC equipment [6, 49].

Due to its high efficiency and low capital costs, active TES systems are often the more commonly used storage method for load shifting applications. However, electrical storage allows a larger portion of the load to be shifted than TES. Even though the production of chilled water is the major source of power consumption in HVAC central plants, other auxiliary equipment, such as fans in air handler units and pumps to move water in pipes, also consume power when delivering the cooling to zones. Active TES cannot shift the power load of these pieces of equipment, which can comprise one-fourth of the total power usage. While battery storage has been expensive previously, prices have been dropping in the past 10 years due to advances in manufacturing and technology development, fueled partially by the popularity of electric vehicles [62].

Furthermore, batteries are able to provide opportunities for revenue generation. Even if they are not used for load shifting of HVAC systems, they can be used generate income by selling the power to the grid during peak periods and buying when prices are cheaper during the off-peak periods [19]. They can also be used to generate revenues by participating in electricity market programs such as frequency regulation [17].

1.2.1 Peak Demand Charge

Load shifting is often not as simple as running HVAC equipment at full capacity to charge active TES tanks and batteries to full capacity during the off-peak periods and not purchasing power during the periods when the prices are higher. In addition to the time-varying usage prices, power companies often levy a *peak demand charge* based on the maximum rate of electricity usage at any moment during the month [13, 52]. These peak demand charges constitute a substantial portion of the total electricity bill, hence they must be considered to obtain economically optimal performance. There is tradeoff between load shifting to decrease usage costs and increasing the peak power load which results in a higher peak demand charge, hence optimization is required for this application. While the peak demand charge has been included in some previous studies, there is little awareness in the HVAC community that including peak constraints may lead to undesirable system behavior [89] if not handled appropriately. Once a new peak has been reached, the controller can operate up to this peak value without penalty; the principle of optimality is no longer valid in such problems since the problem statement changes after achieving a new peak. Therefore, it is not trivial to develop an economically optimal method of HVAC control.

An economically optimal control system needs to be designed for the implementation of these strategies. The research literature reflects a diversity of alternative proposed approaches to HVAC control to augment or to replace simple PID and on/off controllers including genetic algorithms [35, 61], artificial neural networks [38, 45], and fuzzy logic based control systems [43, 91]. Several reviews of these approaches have appeared [2, 18, 28, 37, 90, 97].

However, the focus of this work is primarily on MPC-based control schemes. The reason being that numerous studies have shown that MPC outperforms existing

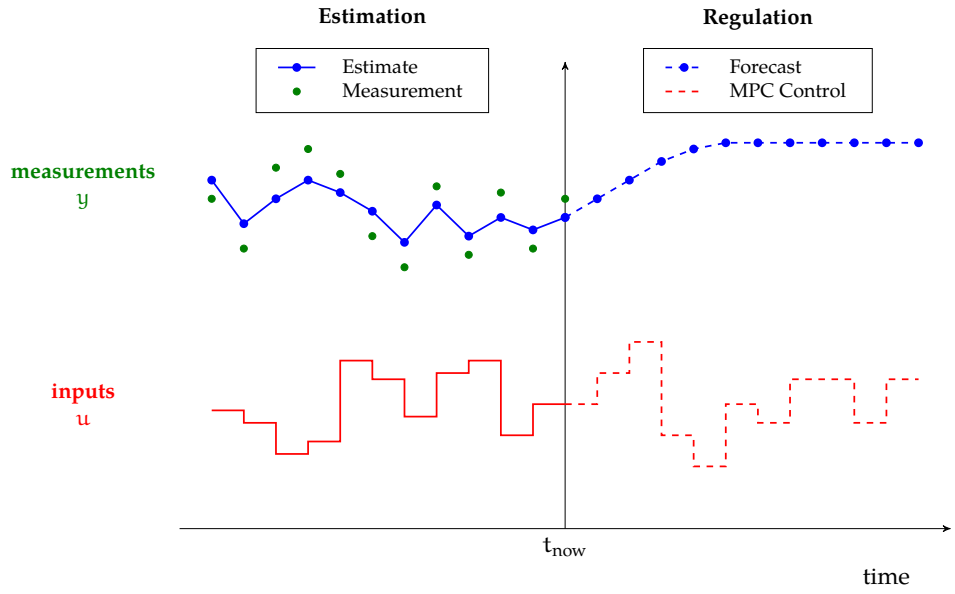

Figure 1.4: Model predictive control paradigm

methods used to control HVAC systems as well as the alternate methods proposed [4, 7, 26, 52, 58].

1.3 Model Predictive Control

As the name implies, *model predictive* control uses a model of the system to make predictions of how the system responds when a certain sequence of inputs is injected, shown in Figure 1.4 [83]. At each step, the controller solves an optimization problem by minimizing a quantity of interest over the forecasting period, or *control horizon*, subject to the system model and a set of constraints on the inputs and outputs. The first control action from the solution of this regulation optimization problem is injected into the system.

In addition to the regulation side of MPC, there is also an estimation part of the MPC problem, depicted on the left side in Figure 1.4. Predictions from the regulation are not always perfect and there can be unmeasured disturbances present in the sys-

tem so measurements of the outputs are used to account for these differences. After each measurement is obtained, a state estimation problem is solved to determine what the current state of the system is using past input and output data as well as the system model. This value of the state estimate is then used by the regulation problem as the starting point for the forecast at the next control time step, and the process is repeated all over.

The formulation allows MPC to handle complex systems with many inputs and outputs with hard and soft constraints. This receding horizon control has had great success in the chemical and petroleum industries over the past two decades [78]. Often, the objective in the regulation problem is tracking error since a common control objective is to drive the system to a setpoint.

For HVAC systems however, the aim is to minimize the total operating cost. Maintaining the system is a fixed setpoint is not necessary; fluctuations in temperature are acceptable provided they are not large. In the last few years, advancements have been made with the development of *economic MPC*, which uses an economic objective function in place of the setpoint tracking one in conventional MPC [3, 21]. Studies have shown that economic MPC enables shifting of the energy load from peak hours to off-peak hours through thermal energy storage (TES) [4, 7, 26, 52, 58, 64]. MPC is well-suited for the control of HVAC applications due to its ability to forecast and manage the various method of energy storage in complex multiple input-multiple output systems. Although researchers in the HVAC community have suggested model predictive control as a promising solution for minimizing energy costs of commercial buildings, the HVAC industry has not yet widely adopted economically optimal control systems.

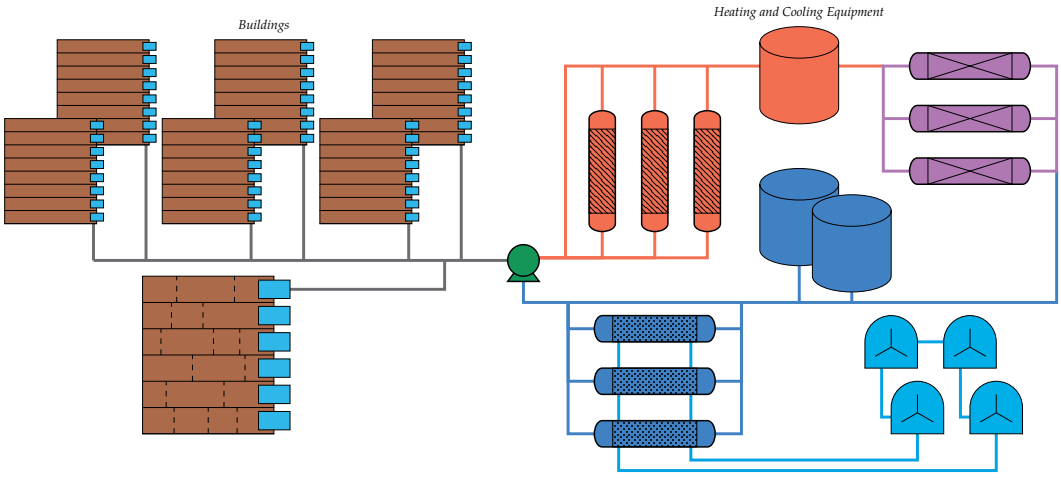

Figure 1.5: Large-scale commercial application with many zones across a campus of buildings and many pieces of HVAC equipment

1.4 Large-Scale Applications

One key reason is that large-scale commercial implementations, such as university campuses and industrial complexes, may contain hundreds of buildings and thousands of zones, as shown in Figure 1.5. A *zone*, in this context, is defined to be a region within a building with a dedicated temperature controller. The MPC controller must determine temperature setpoints for each zone as well as the equipment operation schedule to meet the load required by the buildings. The scheduling of HVAC equipment requires on/off decisions to be made, which introduces discrete variables to the optimization. Solving a single optimization with both continuous and discrete decisions for such large applications over a control horizon on the order of a few days is not practical in real-time, and such control systems are also more difficult to maintain. As a result, decomposing the centralized problem into smaller subproblems is necessary to alleviate these issues and work towards an implementable architecture.

Distributed MPC breaks up the large centralized problem into smaller opti-

mization problems, as discussed in [80, Chapter 6] and [12, 82]. Several iterative methods have been proposed for the buildings problem [11, 20, 46, 88]. However, due to the limitations placed by communication protocols in HVAC systems [59], iterative approaches are not desirable. Some of the proposed methods may involve many exchanges and iterations before converging to the solution. Moreover, peak demand charges cause the objective function to be coupled and decompositions in the presence of peak demand charges have not been addressed. Several viable decompositions have been proposed, but a consensus on which may become the gold standard has not been reached. Developing a case study can provide a standardized system against which to benchmark performance when evaluating decompositions. In this work, the aim is to address these gaps.

1.5 Research Objectives

The greater goal is to design an industrially implementable, economically optimal method of HVAC control. Using energy more wisely not only reduces heating and cooling costs for commercial buildings throughout the U.S. and world, but it also reduces primary energy consumption for power companies since they can operate more efficiently at more nearly constant load, which can help to create a more sustainable future.

In order to help achieve this ambitious long-term vision, the research objectives of this dissertation are:

- to propose a decomposition control strategy for the large-scale commercial applications that addresses the inefficiencies of previous approaches
- to demonstrate that the proposed control architecture can be implemented for industrial systems

- to extend the framework to address other types of HVAC systems including those with embedded batteries
- to release a case study for the HVAC research community to assess alternative control strategies
- to develop nominal stability theory for the zone tracking MPC problem

The remainder of the dissertation is organized as follows. In Chapter 2, the HVAC system and control problem are defined. In Chapter 3, a hierarchical decomposition is proposed for this problem and its benefits are discussed. In Chapter 4, the development of building models, a key step in the implementation of such control systems, is addressed. In Chapter 5, the framework is extended to include distributed battery storage and address variable flow refrigerant systems. In Chapter 6, a large-scale case study loosely based on the Stanford University central HVAC plant is developed and made publicly available for HVAC researchers to provide a basis for evaluating other decomposition strategies. In Chapter 7, an alternative method for tracking zone regions is outlined, and its advantages over conventional setpoint tracking is demonstrated via simulation examples. In Chapter 8, the primary findings are summarized and future directions are considered.

2 PROBLEM DEFINITION

The overall HVAC system can be separated into *airside* and *waterside* systems as denoted in Figure 2.1. The airside system consists of all the zones in the campus buildings along with the temperature regulation equipment such as air-handler units (AHUs). Passive TES exists in the form of building in the airside system. The waterside system consists of the central plant equipment that is used to meet the load generated by the airside system and includes active TES tanks. Electrical storage is considered in Chapter 5.

A supervisory control system is required for each system [69]. Airside controllers determine the cooling duty that must be supplied to AHUs in order for AHUs to deliver supply air at appropriate temperatures to keep the building occupants comfortable. The waterside controller determines how best to operate the chillers and thermal energy storage units to meet a given cooling load. Models are developed for both systems.

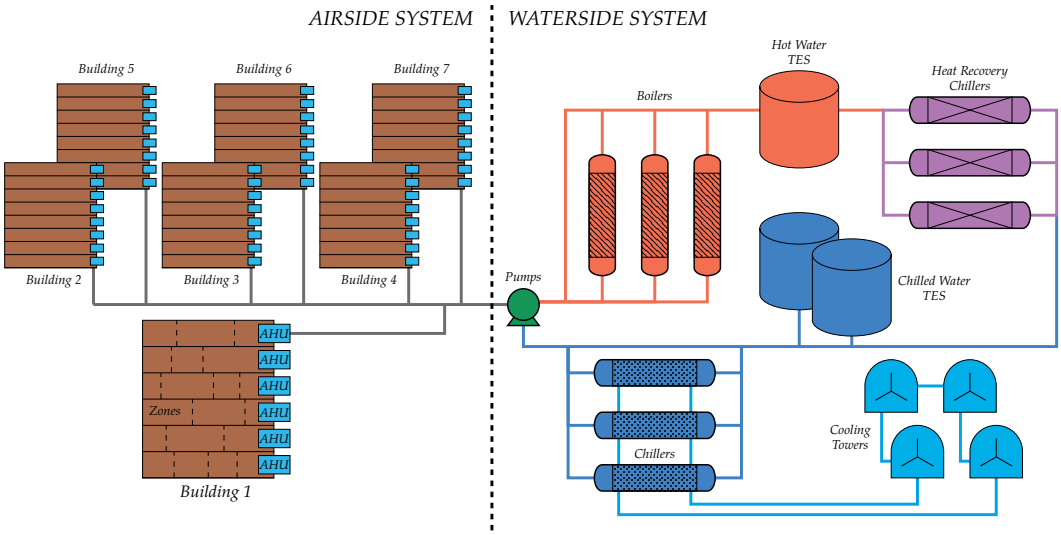

Figure 2.1: Large-scale commercial HVAC application with airside (left) and waterside (right) systems

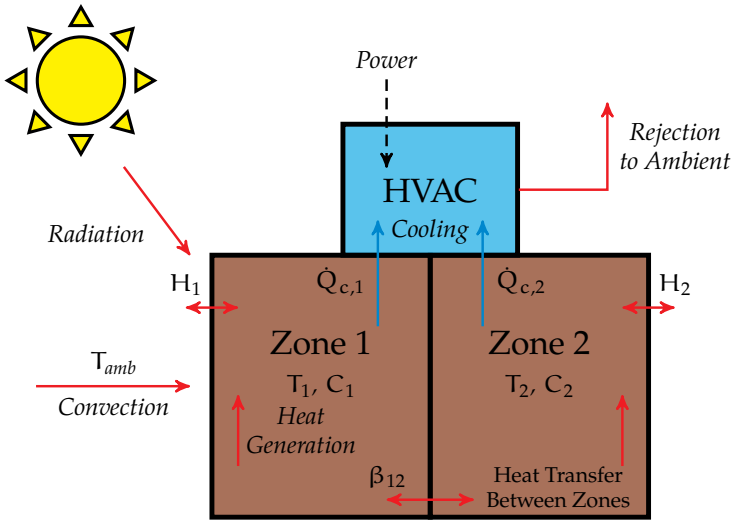

Figure 2.2: Modes of heat transfer present in airside system.

2.1 Airside System

The flow of heat transfer for the airside system is shown in Figure 2.2. The buildings in the campus are divided into zones. Each zone has its independent local temperature regulation controller (e.g., PI thermostat). Each zone may receive heat loads due to radiation, convection with the ambient, and internal heat generation from occupants and lighting. Adjacent zones in the same building also exchange heat with one another, a source of coupling. These heat loads are rejected back to the ambient by the HVAC system (waterside equipment).

2.1.1 Models

A model of the airside system is required in order to implement MPC. The model must describe the temperature dynamics of the building. Small buildings or single zones within a larger building can be modeled with a single temperature and simple linear dynamics [55], but massive commercial buildings are rarely at a uniform temperature. To increase the fidelity of the system model, different temperature

variables can be assigned to different zones of the building, and different variables can be used to model temperatures of air, inner walls, and outer walls [58, 95].

Since these temperature dynamics are often nonlinear and include large time delays, highly accurate models are too complex for online use. Using complex models also provides diminishing returns in terms of the closed-loop performance [76]. Therefore, knowing how to reduce the order and complexity of the model while maintaining sufficient accuracy for effective control, is key in achieving widespread implementation of MPC. Touretzky and Baldea [95] provide a discussion on model reduction based on the dominant time scales of these systems.

Temperature Dynamics. In this work, a lumped approach is used to represent the zone temperature as a single variable. If this approach is not sufficiently accurate, a higher order model may be used instead, however the MPC optimization becomes more difficult and the added complexity may not lead to higher savings [76]. An energy balance for each zone i is expressed to relate heat transfer with its temperature T_i as given by

$$C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{load,i} \tag{2.1}$$

in which C_i is the thermal capacitance, T_i is the zone temperature, H_i is the scaled heat transfer coefficient with the ambient, T_a is the ambient temperature, β_{ij} characterizes the degree of coupling between zones i and j, $\dot{Q}_{c,i}$ is the cooling rate from the HVAC system, and $\dot{Q}_{load,i}$ is the external load on the zone due to other factors such as occupancy and solar radiation. If zones i and j are not adjacent, then $\beta_{ij} = 0$.

Regulatory Control. In this framework, the MPC controller serves as the supervisory control layer that sits above the existing regulation control layer present in

buildings. Hence, the cooling rate $\dot{Q}_{c,i}$ cannot be manipulated directly in practice. The control actions for the MPC layer are temperature setpoints to zone thermostats, which then induce a cooling duty.

Typically, regulatory control is not modeled in supervisory controllers. However, the regulatory control model is required in HVAC systems to relate the temperature setpoint T_{sp} to the cooling duty $\dot{Q}_c$ since that is the quantity optimized by the MPC layer. Responses to changes in temperature setpoint tend to be slow in large buildings so a time-scale separation and $T = T_{sp}$ cannot be assumed.

To represent the zone cooling duty as a function of the setpoint and temperature, two relationships are modeled: (i) the thermostat's control law to determine its output as a function of zone temperature and temperature setpoint, and (ii) an energy balance relating the thermostat output to the actual cooling duty. Assuming an ideal proportional-integral (PI) control law for the thermostat and a linear relationship for the energy balance, the resulting model is

$$\begin{aligned} \dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right] \\ \varepsilon_i &= T_{sp,i} - T_i \end{aligned} \tag{2.2}$$

in which ε_i is the tracking error and $\dot{Q}_{ss,i}, K_{c,i}, \tau_{I,i}$ are the PI controller parameters. Analogous to the temperature case, higher order models can be used for greater fidelity if needed. The approach used to generate model parameters for temperature dynamics and regulatory control is described in Section 4.1.

Disturbance forecasts. Control systems also require adequate disturbance models and forecasts. The disturbances for HVAC systems include occupancy loads, weather, electricity pricing, etc. The primary disturbance is ambient temperature. Stochastic modeling can be used in conjunction with weather forecast data to formulate accurate

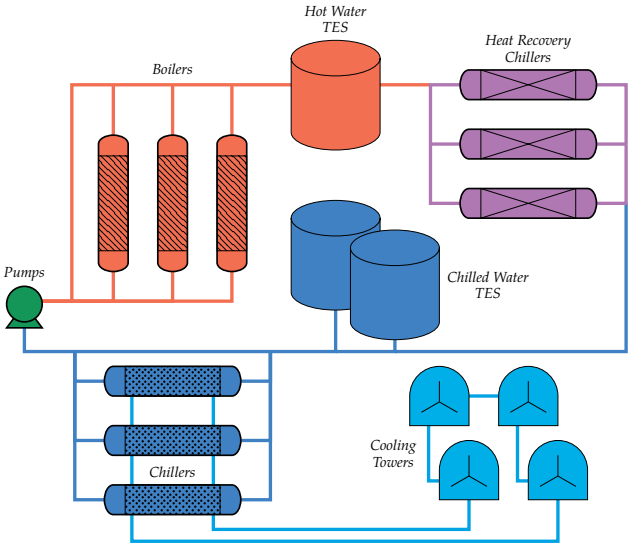

Figure 2.3: Waterside equipment at the central HVAC plant

weather profiles for the controller [65, 101]. Likewise, stochastic modeling can also be used for forecasting electricity market prices. Additionally, loads placed on the HVAC system by building usage are usually predictable with historical data on meeting schedules and typical working hours for employees. Solar radiation can also be included in the model [30]. For this work, it is assumed that these forecasts are already available.

With these disturbance forecasts and models of the airside system in hand, predictions of the system response can be made and the HVAC loads required to satisfy constraints can be estimated.

2.2 Waterside System

The equipment present in the waterside system is highlighted by Figure 2.3. Chillers consume electricity to produce chilled water for meeting the cooling load from the airside system. Associated with the chillers are cooling towers that reject the heat

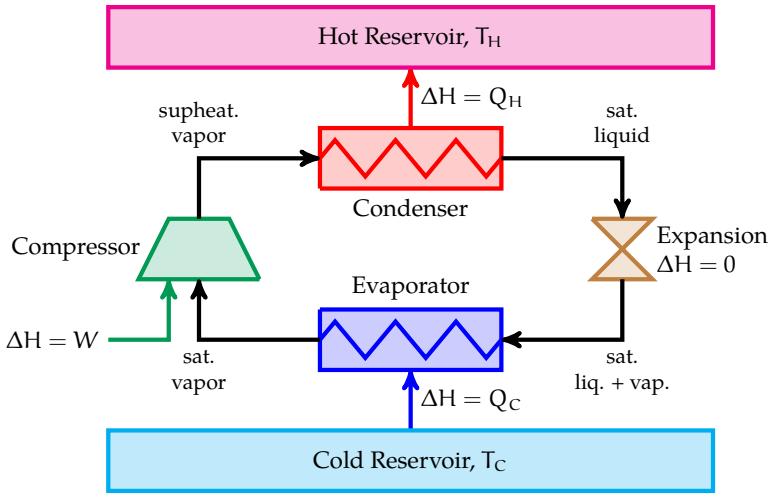

Figure 2.4: Vapor compression cycle.

back to the ambient. Boilers use natural gas to produce hot water for meeting the heating load. Heat recovery chillers (HRCs) use electricity to transfer waste heat from the chilled water return stream to preheat the hot water supply stream. This recycling of heat in the system by HRCs significantly increases the efficiency of the HVAC system. Insulated TES tanks are used for active storage of hot and chilled water.

2.2.1 Models

Similarly, as in the airside case, models are needed for the waterside system to implement MPC. For equipment, power consumption is predicted based on their usage rate. For tanks, storage level dynamics are used to keep track of the inventory.

Equipment Models. Waterside models are specific to the equipment type. For example, the vapor compression cycle for refrigerant inside a chiller is explained in Figure 2.4. In order to produce chilled water, the chiller consumes electricity at the compressor to provide cooling at the evaporator. Its efficiency at converting work

into cooling is summarized using coefficient of performance (COP), which is given by

$$\mathrm{COP} = \frac{Q_C}{W_{\mathrm{net}}} \tag{2.3}$$

The net work is used since additional power may be consumed outside of the compressor itself in transporting fluids. Modeling the detailed thermodynamics is not necessary at the supervisory level. The transient time scale (minutes) is less than that of the scheduling time scale (hours). As such, the equipment are assumed to be at steady state in the supervisory layer and regulatory waterside controller models are not needed. Static empirical models of the form

$$\dot{W}_{\mathrm{equipment}} = f(\dot{Q}_{\mathrm{equipment}}, T_a, ...) \tag{2.4}$$

are developed by fitting equipment usage data for the MPC controller. In addition to the equipment load, power consumption may depend on secondary variables such as ambient conditions. Nonlinear models of this type are developed for all of the waterside equipment, including chillers, pumps, and cooling towers. These models may also be available from HVAC vendors based on the equipment type and model.

Storage Dynamics. A first-order linear model with decay is used for active TES. The storage level, s, evolves according to

$$\frac{ds}{dt} = -\sigma s + \eta_{\mathrm{storage}} \dot{Q}_{\mathrm{storage}} \tag{2.5}$$

in which σ represents the decay term, η_{storage} represents the insulation efficiency, and $\dot{Q}_{\mathrm{storage}}$ represent the charging (or discharging when negative) rate. Typically, active storage tanks are insulated well and their cooling capacity does not decay significantly with time. Approximately 99% of the cooling capacity can be maintained

from one hour to the next.

Discrete Decisions. The waterside equipment introduces discrete variables to the problem. For instance, pumps can be only turned on or off (represented as binary variables, 0 or 1, in the model). We can have only entire chillers operating (represented as integer variables in our model). However, valves and fluid flow rates can take on continuous values. Therefore, the waterside system can be expressed as a mixed-integer programming (MIP) scheduling model. These MIP scheduling models can be written in state-space form [94], which is typically used for MPC applications, allowing for easier integration with the airside system. Note that the discrete and continuous nature of the variable is restricted to actuators only. Full hybrid dynamic models to capture the interactions of continuous and discrete dynamics are not required [27]. This feature restricts the input feasible space.

2.3 Control Problem

With this background, the problem statement is formulated. The aim is to find a solution to the overall control problem: given a forecast of electricity prices, ambient conditions, and disturbance loads, what is the *best* way to operate the HVAC system to minimize the total operating cost. In doing so, several decisions must be made:

- What are the temperature setpoints sent to the zone thermostats?
- What is load generated by the airside system?
- When is storage being charged and when is it discharged?
- When are the waterside equipment turned on and at which loads are they running when operating?

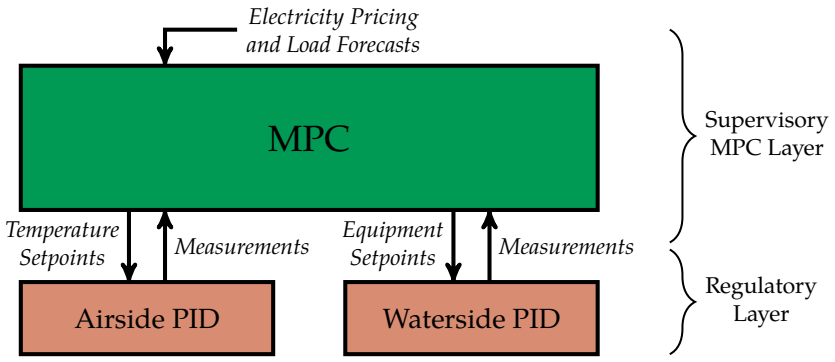

Figure 2.5: Combined MPC structure for the airside and waterside systems

While making these decisions, the following constraints must be respected:

- The zone temperatures must be kept in a predefined region to maintain occupant comfort.
- Air-handler units, waterside equipment, and storage tank have physical capacity limits according to their size.
- Chillers and other large equipment cannot be switched on and off too rapidly.

2.4 Combined MPC

One way to design a control system for this problem is to formulate and to solve a single monolithic optimization problem. This approach, shown in Figure 2.5, is denoted *combined MPC* since the airside and waterside systems are combined into one problem.

The MPC controller receives pricing and load forecasts. It uses detailed models of all airside zones and all water equipment to solve a large optimization problem which minimizes the operating cost. The solution of this problem is sent to the regulatory layer in the form of setpoints. Measurements are sent back to the combined MPC

problem as feedback so the forecasts and disturbance estimates can be update for the next controller time step.

This combined MPC approach may work well for small systems with few zones and few pieces of equipment. However, campus-wide implementations may contain hundreds of buildings and thousands of air handler regions each with tens of zones, as discussed in Chapter 1. A single centralized control system for these applications is impractical and undesirable. In such cases, distributed MPC can be employed [93].

2.5 Distributed MPC Approaches

Special cases of distributed MPC have been studied in the HVAC literature [59, 77]. An iterative approach can be used, whereby the controllers take turns to decrease the global objective by manipulating only their own variables [88]. Although computations can be overlapped, this approach can be slower than solving a centralized problem because of the number of information exchanges that must occur and the slower rate of convergence of the block-coordinate optimization approach. Another iterative approach uses a hierarchical decomposition of the problem into a zone layer (which handles local variables) and a coordination layer (which decides how to split the total power allocation among the zones) [46]. An explicit MPC solution for a horizon length of one period has also been formulated [42], but much longer horizons are needed to take advantage of load shifting. Distributed MPC has also been applied to the related problem of smart power grids [14].

As mentioned above, building temperature control dynamics are highly nonlinear, often with large time delays. Under certain circumstances, the optimization problem may even have multiple local optima, which can correspond to different operating modes or different permutations of heating and cooling the various zones

[40]. These multiple optima can result from the non-convexity and degeneracy of the optimization problem. Several iterative and decomposition techniques have been discussed in the literature for solving the optimization problem at each MPC step in a distributed manner. These include sequential linear programming [64], sequential quadratic programming [53], piecewise linear approximation of nonlinear functions [48], the fast gradient method [56], Benders' decomposition [60], and a bundle method [46, 47]. However, these methods are not universally applicable for every building.

To have a broader impact, the goal is to design a control system that minimizes the information exchanges between controllers so that it can be widely implemented. Additionally, the peak demand charge is a major source of coupling amongst zones even in separate buildings. The decomposition of the centralized problem needs to consider this fact.

2.6 Design Features

In Chapter 3, an economic MPC-based architecture for large-scale building temperature regulation is proposed to address these aforementioned key design issues.

- The overall combined problem is decomposed into smaller subproblems for commercial applications.
- Computational complexity is reduced by using aggregate system models.
- Demand response is addressed in the peak demand charges for the airside and waterside systems.
- The communication burden is reduced by using a non-iterative method.

- The formulation is made general enough to handle a large class of HVAC systems.

3 DECOMPOSITION

In this chapter, we propose a decomposition of the combined problem to address the key design features outlined in Chapter 2. The subproblem formulations are presented and simulations results are used to demonstrate the energy cost savings that can be achieved with this control system.

3.1 Hierarchical Decomposition

A hierarchical decomposition for HVAC control is proposed as shown in Figure 3.1. The MPC layer is divided into high- and low-level problems. The high-level problem performs a system-wide economic application using pricing and disturbance forecasts as well as aggregate models of the system. It serves as the coordinator in the hierarchy. The low-level problem is split into airside and waterside following industrial convention. Using the load allocations computed from the high-level problem, the low-level problems perform further optimization using more detailed models. The formulations of these problems are discussed in more detail later in

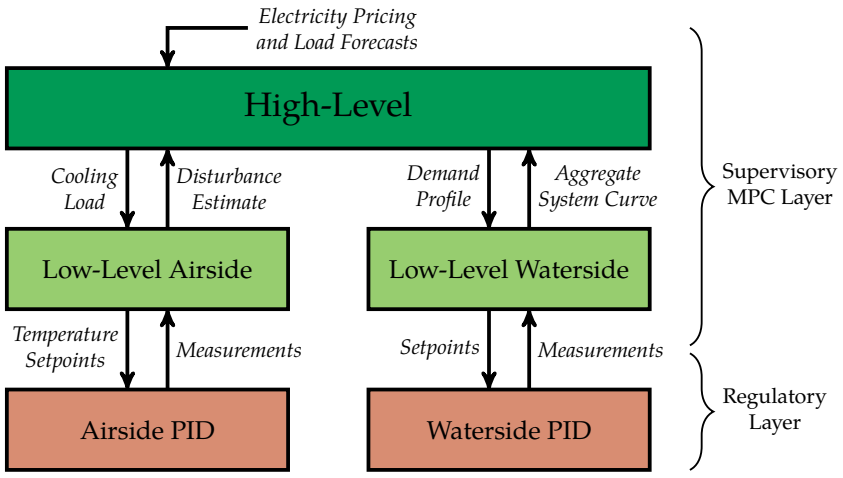

Figure 3.1: Hierarchical decomposition for control of airside and waterside systems

this chapter. The setpoint decisions computed from these low-level problem are implemented in the physical systems.

As before, the MPC layer lies above the regulatory layer, which consists of the building automation system (BAS) that serves as an interface to the PI or PID controllers for the airside and waterside systems. The design is expected to work with any existing BAS, allowing one to retrofit an existing HVAC control architecture with optimization. Only setpoint and measurement communication between the layers is required. Measurements from the BAS are sent back to the MPC layer and used as feedback to update forecasts for the next controller execution.

3.2 Decomposition for Large Systems

For large-scale applications with hundreds or thousands of zones, it may not be feasible to optimize over all zones simultaneously in a single low-level airside optimization problem. For such applications, we further decompose the problem.

First, the entire collection of zones is divided and grouped into a series of *subsystems*. For example, the grouping can be done by building. The first subsystem consists of all zones in Building 1, the second subsystem consists of all zones in Building 2, and so on. While it is not a requirement to group zones into subsystems by building, doing so results in subsystems that are exactly decoupled since buildings do not directly exchange heat with one another. Grouping zones in a manner where they exchange heat with zones in other subsystems due to sharing a physical boundary leads to coupled subsystems and the strength of the coupling may affect the performance of the decomposition. For this reason, we choose to make each subsystem represent a separate campus building in the discussion that follows.

Once this classification is performed, the low-level airside problem can be decomposed into separate MPC controllers with each controlling a subsystem of the

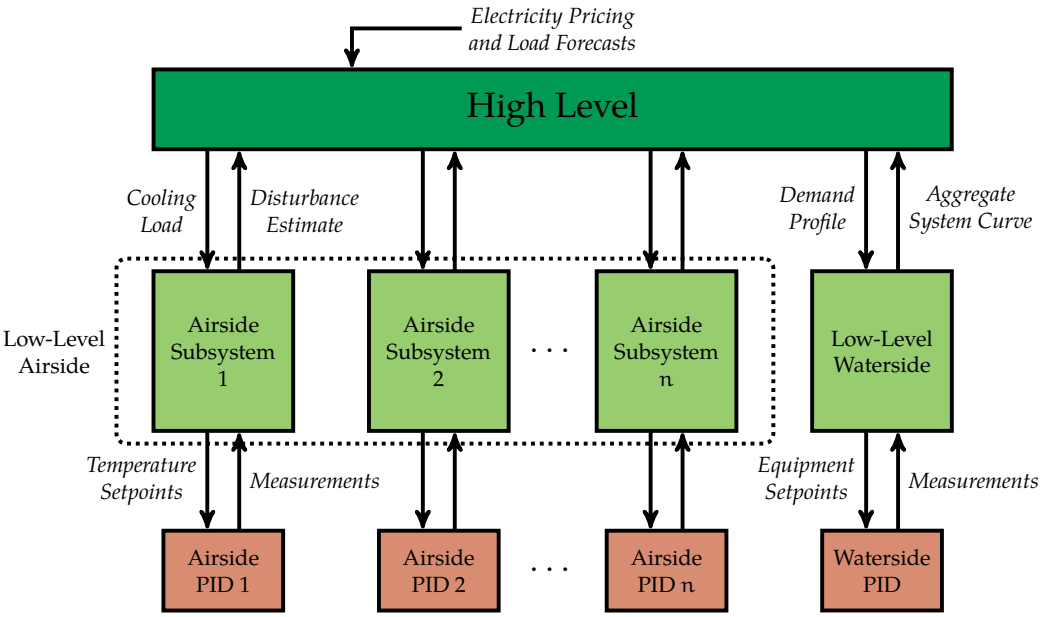

Figure 3.2: Hierarchical decomposition with distributed low-level airside subsystems

entire airside system as shown in Figure 3.2. The information exchanges between the layers are identical for all of the distributed airside controllers as described above. The high-level problem contains an aggregate model of each airside subsystem and allocates a cooling load to each subsystem. The subsystem MPC controllers contain more detailed zone-level models. The choice of dividing the airside system into separate subsystems is made for computational reasons so that the subsystem MPC problems can be solved within a few minutes.

In this decomposition, the high-level problem determines how much cooling load to allocate to each building subsystem to optimize total operational cost. For the system model, the high-level uses an active TES model as well as aggregate airside and waterside models to reduce computational complexity. The results of this computation are sent down the low-level airside and waterside controllers.

Each building subsystem has a separate low-level airside controller that computes the zone temperature setpoints that minimize energy usage while maintaining the

zone temperatures in a prespecified comfort region and not exceeding the cooling load from the high-level problem.

The low-level waterside problem determines the equipment usage that minimizes cost while meeting the load from the high-level problem; its decision variables include equipment operation, thermal loads for chillers, flows for pumps, setpoints for other auxiliary equipment, and storage utilization. Due to the discrete decisions that are made, the waterside problem is formulated as a mixed-integer linear program (MILP). Previously, heuristics have been used by operators at central HVAC plants to make these scheduling decisions.

The setpoints from the low-level problems are dispatched to the regulatory controllers of the BAS. Feedback is employed to communicate measurements (e.g., temperature, power, equipment state) from the regulatory layer to the low-level MPC layer. Then, updated aggregate models and disturbance forecasts for both airside and waterside systems are communicated to the high-level problem for the next MPC execution step.

3.2.1 Advantages of Proposed Structure

The proposed control decomposition has several advantages over previous methods. The two primary sources of coupling (peak demand charge and total equipment capacity constraint) are addressed in the high-level problem, which acts as a coordinator. The high-level problem uses aggregate airside and waterside models to reduce computational burden. Subsystems in the low-level airside problem are decoupled by construction. Hence, iterations between the building subsystems are not required, greatly reducing the internal communication burden. Most building temperature regulation methods do not consider detailed models of chiller plant equipment or integer decision variables, which decreases the fidelity of energy cost

calculations. Including integer variables in the waterside problem allows the optimizer to determine the equipment operation schedule rather than using heuristics. The decomposition allows the architecture to scale well to handle large systems since the subproblems can be solved in real-time for online deployment. The general framework can handle various types of applications: central chiller plants, air handler units (AHUs), roof top units (RTUs), variable refrigerant flow (VRF) systems, and any BAS that can receive temperature setpoints and send measurements. It can also handle both active and passive TES. Next, the high-level and low-level airside problem formulations are discussed.

3.2.2 High-Level Problem

As depicted in Figure 3.2, the input data for the high-level problem includes electricity pricing, weather forecast, building usage model, and disturbance estimates from low-level problem. The economic cost objective includes two components, demand charge and usage charges, since most pricing structures include a peak demand charge on the maximum power usage during the month in addition to the time varying electricity costs based on usage, as discussed in Chapter 1. The decision variables are building temperatures, storage utilization, and thermal loads for waterside system. Dynamics models for the building temperature and active TES are used to make predictions. The average building temperature models have a similar structure to the zone heat transfer models provided in Equation 6.1. The aggregate waterside system model is represented by the coefficient of performance (COP) in the high-level objective. Constraints include comfort bounds on the air temperature and bounds on the cooling duty. The comfort zone constraints are implemented as *soft constraints* and their violations are penalized in the objective. Mathematically, the high-level

MPC problem is formulated as

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u},\dot{Q}_{peak}} \quad & \sum_{k=0}^{N-1} c_k \eta \dot{Q}_{HVAC,k} \Delta + c_{peak} \eta \dot{Q}_{peak} \\
\text{s.t.} \quad & C_b \frac{dT_b}{dt} = -H_b (T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{other,b} \\
& \frac{ds}{dt} = -\sigma s + \dot{Q}_{storage} \\
& \dot{Q}_{HVAC,k} = \sum_b \dot{Q}_{c,b,k} + \dot{Q}_{storage,k} \\
& 0 \leqslant \dot{Q}_{HVAC,k} \leqslant \dot{Q}_{peak} \\
& \dot{Q}_{peak,past} \leqslant \dot{Q}_{peak} \leqslant \dot{Q}_{HVAC,max} \\
& 0 \leqslant s_k \leqslant s_{max} \\
& T_{min} \leqslant T_b \leqslant T_{max}
\end{aligned}
\tag{3.1}
$$

Since the models are linear, this optimization is a linear program (LP). The representation of the high-level MPC problem as standard form LP is presented in Appendix A.

As a MPC problem, the states for the high-level controller are the building temperatures and active storage tank level. The inputs are the cooling duties to each building and the amount of storage charged or discharged. The disturbances are the ambient temperature and external loads placed on each building. A forecast of this disturbance vector is obtained and provided to the optimization problem at each MPC step. In the simulation study, the states are assumed to be measured directly. In industrial applications, state estimation is required to reconstruct the states from the measurements. A rate of change penalty can also be added to the inputs if needed.

Model Aggregation. To derive building models from zone models for the high-level problem, model aggregation is employed. The average building temperature is defined by summing over all zones in that building using

$$T_b = \frac{\sum_i C_i T_i}{\sum_i C_i} \tag{3.2}$$

Using this definition, building parameters for the high-level airside model are computed by summing parameters from the low-level airside models using

$$C_b = \sum_i C_i \qquad H_b = \sum_i H_i \qquad \dot{Q}_{other,b} = \sum_i \dot{Q}_{other,i} \tag{3.3}$$

The derivation for this model aggregation method is given in Appendix A. Model aggregation also helps to reduce the amount of information that needs to be exchanged, thus, addressing the potential practical limitations with existing HVAC communication protocols. For the waterside system, the aggregate model is represented using the COP.

3.2.3 Low-Level Airside Problem

As shown in Figure 3.2, the input data for the low-level airside problem includes the cooling duty allocated from the high-level problem. The objective is to minimize energy usage. The decisions are the zone temperature setpoints for all zones in that subsystem. Constraints include comfort bounds on the zone air temperatures and bounds on the cooling duty obtained from the high-level problem; they are implemented as soft constraints with comfort zone violations penalized more than high-level cooling duty violations. The zone temperature dynamic models are used, as presented in Equation 6.1. The zone air temperature and temperature setpoint to cooling duty model ($\dot{Q}_{c,i} = f(T_i, T_{sp,i})$) is given by Equation A.3.2.

In each building, the low-level airside MPC problem is formulated as

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u}} \quad & Q_{\text{total},N} \\
\text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \\
& \frac{dQ_{\text{total}}}{dt} = \sum_i \dot{Q}_{c,i} \\
& T_{\min} \leqslant T_i \leqslant T_{\max} \\
& \frac{Q_{\text{total},k+1} - Q_{\text{total},k}}{\Delta} \leqslant \dot{Q}_{\text{HighLevel},k} \\
& Q_{\text{total},k+1} - Q_{\text{total},k} \geqslant 0 \\
& \dot{Q}_{c,i} = f(T_i, T_{\text{sp},i})
\end{aligned}
\tag{3.4}
$$

Since the constraints and objective are linear, this optimization is a LP as well. The representation of the low-level airside MPC problem as standard form LP is presented in Appendix A.

Temperature Tracking Another alternative strategy for the low-level airside problem can be to track the average building temperature computed from the high-level problem. In this setup, instead of sending the computed cooling load allocation, the high-level controller sends its computed average building temperature profile down to the low-level airside controllers. The aim for the low-level airside controllers is then to track the average building temperature while maintaining the zone temperatures in the comfort region.

Note that tracking the average building temperature may not result in a unique solution for this case. Consider an example with two zones. The high-level controller decides that the optimal average building temperature should be 21 °C. If the comfort zone is between 20 °C and 22 °C, the low-level airside controller can set zone 1 to

be 20 °C and zone 2 to be 22 °C or vice versa to track the average temperature. It can also set both zones to 21 °C and achieve the same objective. Since there exist nonunique solutions for this problem, a small penalty on the total energy usage is added to pick the "best" of these solutions. Perhaps, it is more cost effective to maintain zone 1 at a higher temperature since it contains many windows and solar radiation places a larger burden on its air handling unit.

To track the temperature computed from the high-level problem, the high-level cooling duty constraint from (3.4) is removed, and the objective of the problem is replaced with

$$\sum_{k=1}^{N} \frac{1}{2} \| T_{b,k} - T_{HighLevel,k} \|^2 + \mu Q_{total,N}$$

in which the average building temperature is defined by Equation (3.2) and μ is a small penalty placed on energy usage. This penalty on energy usage is necessary to ensure a unique solution as discussed above. With this objective, the problem becomes a quadratic program (QP).

Based on simulations, there is not a significant difference in performance between the two objectives of energy minimization and temperature tracking, hence both are suitable for use in the low-level problem [70].

3.2.4 Low-Level Waterside Problem

The total HVAC load from the high-level problem needed to satisfy the airside system is sent to the low-level waterside problem. The waterside controller determines the operation schedule of the equipment to meet the load, including when individual pieces of equipment are turned on and off (discrete decisions) as well as what load they running at (continuous decisions). The objective is minimize operating cost. Demand charge and usage costs are considered. Detailed equipment models are used as well as active TES models. Piecewise linear approximation are made of the

nonlinear equipment performance models so that the problem can be posed as a MILP. The computed solution is then sent to the central HVAC plant equipment. Risbeck et al. [85] provide further details on the mathematical formulation of the low-level waterside problem. In this dissertation, the focus is primarily on the decomposition strategy and formulations of the high-level and low-level airside problems.

3.2.5 Feedback

After the control actions are executed, measurements are taken. Models are also not perfect and unknown disturbances can enter the system. These measurements are fed back to the low-level MPC controllers from the regulatory layer and are used to estimate unmeasured disturbances and update any disturbance forecasts for the next time step. Additionally, as the central HVAC plant operation changes, the aggregate waterside model in the high-level problem can be updated on a slower time scale to account for these changes.

3.3 Simulation Results

To demonstrate the effectiveness of this control architecture, an illustrative example is presented for a modest-sized campus with 50 zones. However, the architecture is scalable and can be easily extended to handle large campuses with hundreds of zones and tens of buildings as shown in Chapter 6.

3.3.1 System Setup

In this example, a campus of 10 buildings, each with 5 zones, is considered. The cooling loads for these 50 zones are provided by a central HVAC plant with active TES. The weather and electricity pricing data used in the simulation are shown in

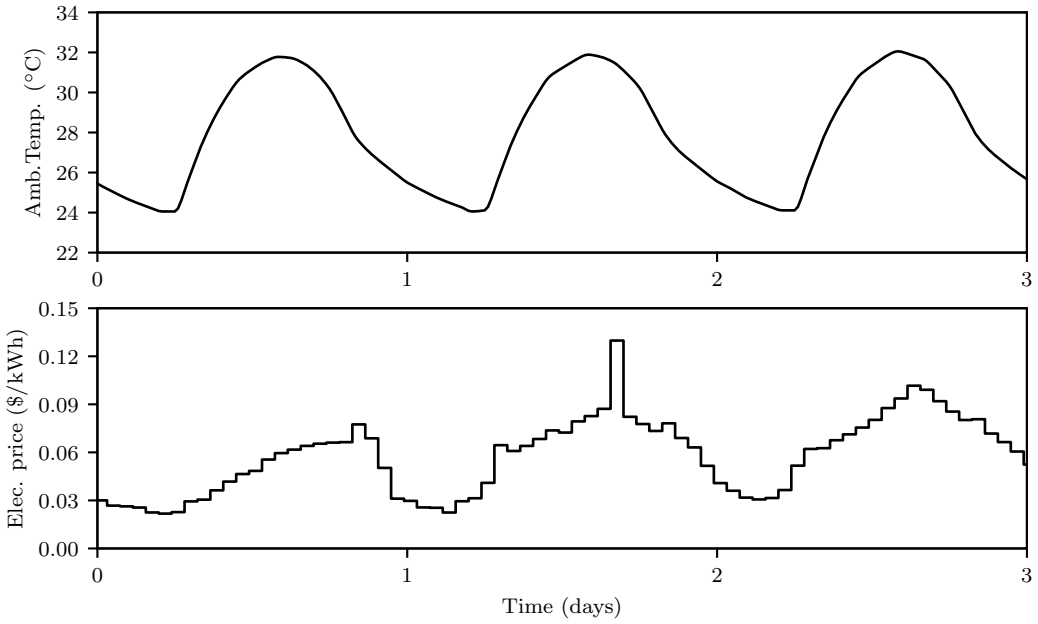

Figure 3.3: Weather and pricing data

Figure 3.3. The weather data were obtained from Atlanta, Georgia for a three-day period in the summer. The time-varying electricity price and peak demand charge data are representative of data from Johnson Controls. The comfort zone for each zone is between 20.5 °C and 22.5 °C. The horizon is 3 days, and timestep is 15 min.

The central cooling plant contains eight conventional chillers, six pumps, ten cooling towers, and a small active TES tank. The minimum and maximum capacities of each chiller are 2.5 MW and 12.5 MW, respectively. The maximum cooling capacity of the storage tank is 100 MWh.

3.3.2 Optimization Results

The optimization problems were solved using Gurobi 7.5.2 via MATLAB R2017b on a desktop machine with 16GB RAM and 2.70GHz Intel® Core™ Processor i5-7500T. Figure 3.4 shows the results of the high-level optimization problem. The top graph

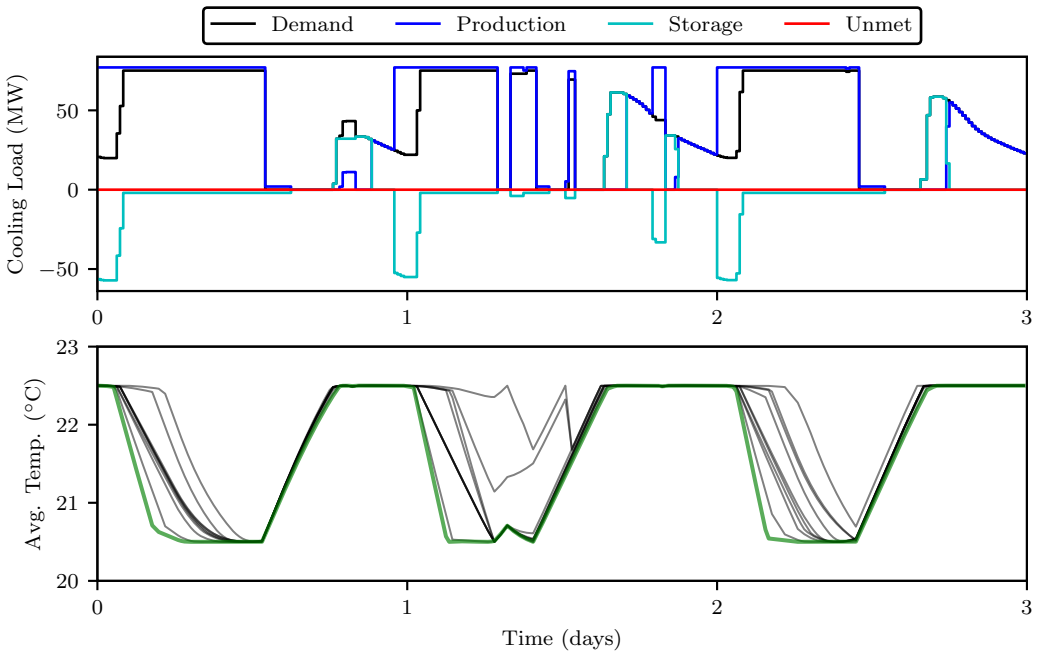

Figure 3.4: Results from the high-level optimization

plots the total campus load and how it is met via equipment production and active storage. The bottom graph shows the trajectory of average building temperatures across the campus. With passive TES available, the economically optimal strategy is to precool the buildings before the peak period when electricity prices are highest. Active storage is also charged during this period. When the prices are high, the storage is discharged so the equipment can be turned off and the buildings are allowed to heat up to the upper bound of the comfort zone. The flatness of the cooling duty profile from the high-level problem is a result of the peak demand charge. Note that the precooling strategy is not always trivial. On the second day, the building temperature profiles are different compared to the other days due to a spike in the electricity price. Since the electricity market can be volatile, optimization is required to determine the minimal cost operating strategy.

Figure 6.3 shows the results obtained by solving the ten low-level airside prob-

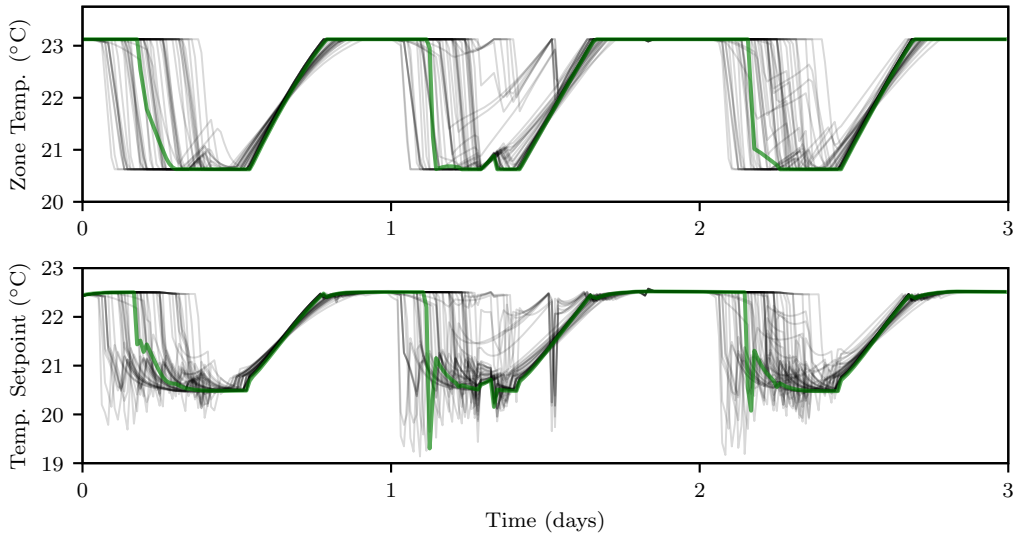

Figure 3.5: Results from the low-level airside optimization

lems. The top graph shows the 50 zone temperatures, and the bottom graph shows the setpoints computed. The low-level airside problem determines the order of precooling based on the dynamics of the individual zones in the building to minimize energy usage. The passive TES is utilized in all three days due to the load allocation from the high-level problem. One interesting feature to note is that the setpoints trajectories often have overshoot. The MPC layer is accounting for the sluggish response dynamics in the buildings by sending a lower setpoint early to control the precooling rate for the zones. Hence, MPC can also be used to improve the temperature response in older buildings with poorly tuned PI controllers.

Figure 3.6 shows the results of the low-level waterside optimization problem. The top graph shows how the demand profile from the high-level problem is met using a combination of chiller production and active storage tank. The bottom graph shows the corresponding equipment operation schedule required to satisfy the load. The shaded region depicts when the equipment is turned on with the approximate

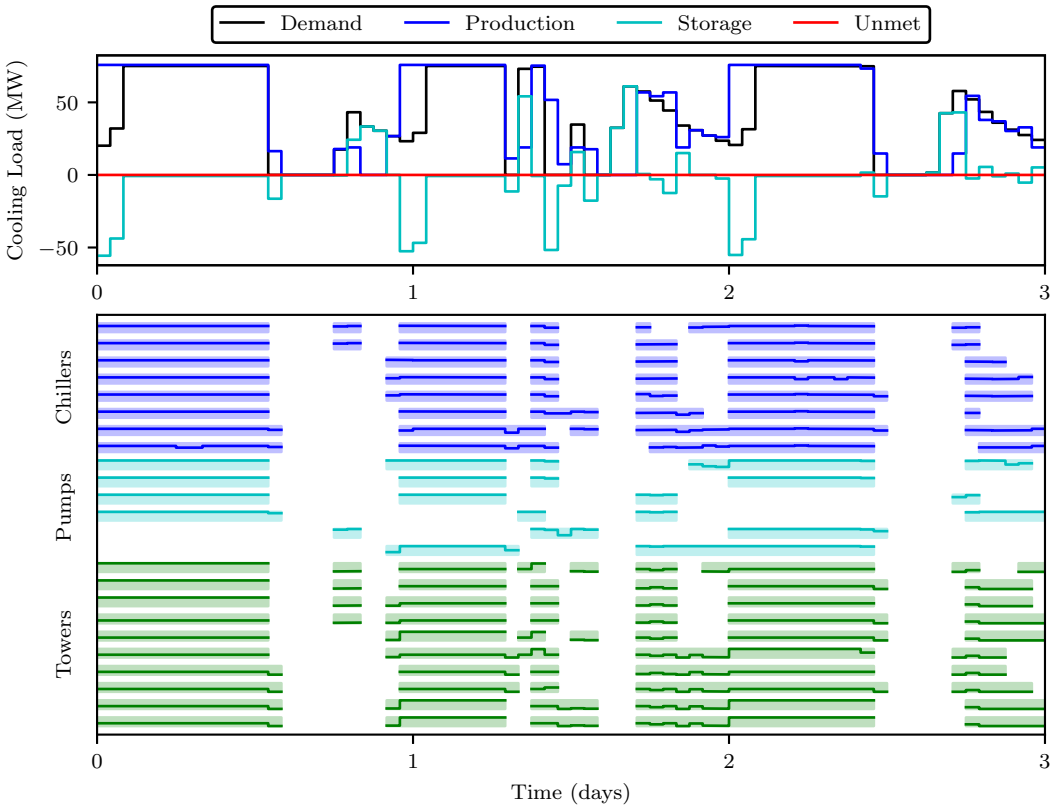

Figure 3.6: Results from the low-level waterside optimization

loading fraction between 0 and 1 depicted by the solid line. The production schedule is similar to the one computed from the high-level problem, but the waterside is able to capture the features caused by the discrete nature of equipment. With active TES available, chillers are run during the night to charge the storage tank and they are turned off or run at lower rates during the peak periods when prices are more expensive.

The operating cost for this policy is $32,402. To determine the value added by optimization and load shifting with this approach, the performance is compared against a baseline. One intuitive baseline is to avoid precooling by employing the minimum energy usage strategy: stay at the upper limit of the comfort region at

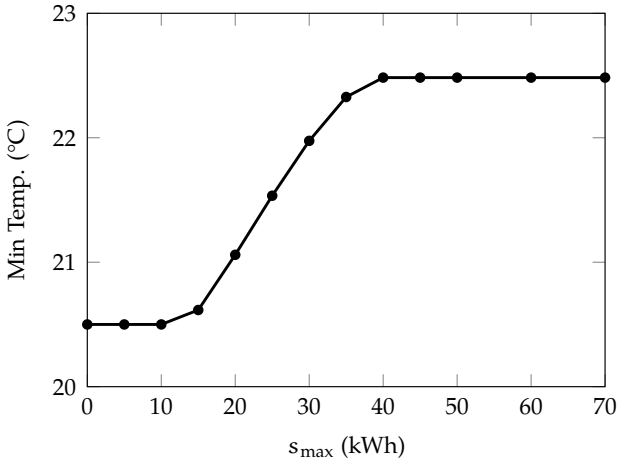

Figure 3.7: Effect of active TES capacity on the usage of passive TES

all times. This baseline uses a setpoint tracking controller. When employing this strategy, the cost is $38,891. Hence, the proposed control system is able to provide savings of 16.7% with a modest amount of energy storage.

3.4 Storage Sizing Study

The capacity of the active TES tank plays a key role in determining the optimal precooling behavior. Figure 3.7 shows the results from a study conducted on one particular system where the size of the active storage was varied and the amount of precooling performed by the control architecture was measured. The minimum temperature across all zones over the prediction horizon in the optimal solution of the low-level problem is plotted as a function of the active TES tank capacity. As the storage size increases, the minimum temperature increases, suggesting that the controller is not precooling the zones as much. After 40 kWh capacity, the controller stops precooling zones altogether for this system.

This study shows that the economically optimal control policy is to stay at the upper bound of the comfort zone at all times if enough active TES is available.

This result occurs due to the high efficiency of active storage compared to passive storage. If the storage tank is large enough, the active TES is able to perform all of the load shifting and additional economic incentive does not remain to precool the buildings since the equipment are turned off already during the peak period. The exact value of storage capacity required to reach this point depends on a variety of factors: disturbance loads placed on the buildings, electricity pricing structures, and physical parameters of the system. However, this optimization framework provides an accurate method of computing this value by varying s_{max} and observing the precooling response; this approach can be useful for storage tank sizing during the design of central HVAC plants.

During these studies, it was concluded that energy cost savings can range from 5% to 40% for typical HVAC applications, depending on the amount of active TES available in the system and the price difference between peak and off-peak periods. With no active TES, significant savings are not possible with inefficient passive TES only. The savings increase as active TES is added to the system and saturate once the entire load has been shifted. The maximum amount of cost savings depend on the difference between peak and off-peak pricing in the electricity market.

3.5 Summary

A hierarchical decomposition has been proposed for large-scale commercial HVAC applications. The decomposition uses a high-level controller with aggregate models and low-level controllers with higher fidelity models. The economic benefits of load shifting using this decomposition strategy have been demonstrated on a sample system. The energy cost savings increase with the amount of TES available to the system and after a critical point, additional savings are not possible since it is not physically possible to shift the load further.

4 IMPLEMENTATION

In order for these control systems to be implemented, several steps are required before deployment can occur. One key step is to develop models for making predictions in the MPC controllers. In this chapter, the methodology for HVAC model development is discussed. Additionally, airside power consumption is explored as an additional opportunity for achieving cost savings via load shifting.

4.1 Motivation for System Identification

Thus far, the estimation and regulation problems in MPC have been discussed in Chapter 1. In the estimation problem, the model of the system G is used along with input u and output y data to estimate the state x of the system. In the regulation problem, the model G and desired output y are used to compute the optimal input u to inject into the system. Both of these problems require knowledge of the system model G.

System identification (SI) is the procedure used to compute the model G from input-output data, as shown in Figure 4.1 [51]. Model identification can performed in the frequency domain with transfer functions or in the time domain. In this work, the time domain to obtain state-space models. The SI problem is formulated as a nonlinear program (NLP) to determine the model parameters θ which minimize the

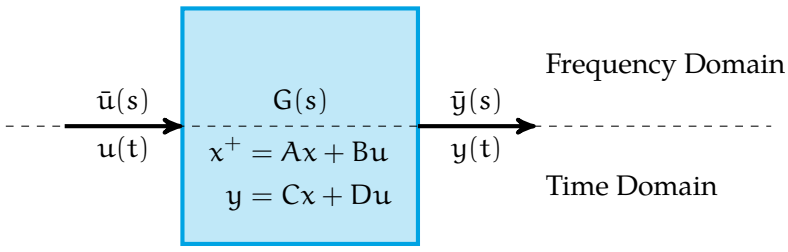

Figure 4.1: System identification

error between model predictions $\hat{y}$ with the measurements y

$$\min_{\theta} \sum_{k} \|y_k - \hat{y}_k(\theta)\|^2$$

This parameter estimation problem can be solved using identification toolboxes, e.g. MATLAB's SI Toolbox.

4.2 Black-Box and Grey-Box Models

There are several choices available in terms of the model form, $\hat{y}(\theta)$. In this section, linear *black-box* and *grey-box* models are addressed.

In a linear black-box model, no knowledge of the internal structure of the system is assumed besides linearity. The u and y data are expected to define the model entirely. The key benefit of this approach is that the SI techniques can be applied to any system without having to supply any preexisting knowledge about the system, which decreases the user input required. However, the disadvantage is that without this knowledge of the internal structure, the model has more unknowns and tends to be less accurate.

Linear grey-box models provide the capability to add domain knowledge to the SI approach, enabling easier, faster, and more accurate model identification. The structure of the model is specified by user beforehand and only a small subset of unknown parameters remain to be identified from data. The data requirements may drop significantly after adding this structural information to a grey-box model. Less data mean less expensive testing required, and more accurate models may result in better control when using model-based control methods like MPC.

In addition to grey-box and black-box models, there are also *white-box* models, in which the entire model is derived from first principles without the need for

Figure 4.2: PRBS input excitation

operational data [75]. While not requiring data is a positive, these models are more difficult to develop since they require expert knowledge of the system, which can increase the labor costs. Additionally, the resulting models derived from partial differential equations (PDEs) can have thousands of states which makes them difficult to optimize and hard to scale up for large applications. Additionally, the added complexity often does not result in a noticeable improvement in closed-loop MPC performance [76]. Hence, white-box models are not considered in this work. The grey-box and black-box approaches are evaluated instead.

4.3 Input Excitation

In order to perform identification, a dataset needs to be generated. Using closed-loop historical operating data is often not recommended since there is a high degree of correlation in the variables which can make it difficult to discern the effect of each input on each output. Sufficient independent excitation in the inputs is needed, which

is performed during identification tests. For this input excitation, one method is to use a pseudo-random, binary signal (PRBS) in which the inputs are switch between two values at pseudo-random points in time. This approach is less disruptive to the output and has a higher frequency content than step testing. Since a large amplitude PRBS can result in frequent saturation of the PID controllers in air handler units, small amplitude PRBS is used instead, as shown in Figure 4.2. Moreover, temperatures must be kept in the comfort zone if the building has occupants during testing period hence the amplitude of the PRBS was chosen to satisfy the comfort bounds. The inputs from Figure 4.2 are injected into the system and the measurements are collected to generate the u and y data for SI.

4.4 Types of Grey-Box Models

There are several choices for the type of grey-box model that can be used. In this section, three popular model forms are presented:

1. Lumped Model
2. Air/Mass Model
3. Air/Shallow/Deep Model

The complexity increases with each model using additional state variables to characterize higher order phenomena and various time scales in zone temperature dynamics. Increased complexity does not always lead to significant improvement in closed-loop performance [76]. Heat transfer in zones can often be captured well enough with just a few states. But these alternative model forms are shown for completeness.

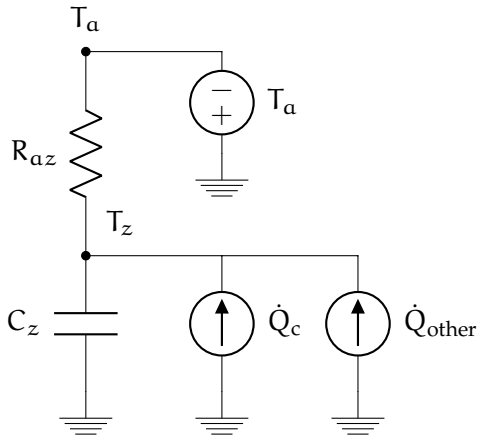

Figure 4.3: Network for zone lumped model where $R_{(\cdot)} = \frac{1}{H_{(\cdot)}}$.

4.4.1 Single Zone Lumped Model

For the single zone lumped model, a single state is used to model the zone temperature. With a PI controller, the model for each zone i is given by

$$\begin{aligned} C_{z,i}\frac{dT_{z,i}}{dt} &= H_{az,i}(T_a - T_{z,i}) - \dot{Q}_{c,i} + \dot{Q}_{other,i} \\ \dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right] \\ \varepsilon_i &= T_{sp,i} - T_{z,i} \end{aligned}$$

in which $C_{(\cdot)}$ is a thermal capacitance and $H_{(\cdot)}$ is a scaled heat transfer coefficient. This lumped model form was used in the identification work presented earlier in this chapter. Often in the HVAC literature, heat transfer models are presented using RC circuit form. A depiction of the corresponding thermal network is shown in Figure 4.3 with $R_{(\cdot)}$ representing heat resistances. While the RC model form is frequently used in electrical engineering works, the heat transfer form with $H_{(\cdot)}$ is the preferred choice in this document.

As shown in Table 4.1, the states are the zone temperatures, the integrals of the

States	$x = \begin{bmatrix} T_{z,i} & \bar{\varepsilon}_i & Q_{total} \end{bmatrix}^T$	$n = 2n_{bz} + 1$
Inputs	$u = \begin{bmatrix} T_{sp,i} \end{bmatrix}^T$	$m = n_{bz}$
Measurements	$y = \begin{bmatrix} T_{z,i} & Q_{total} \end{bmatrix}^T$	$p = n_{bz} + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,i} & \dot{Q}_{ss,i} \end{bmatrix}^T$	$n_d = 1 + 2n_{bz}$

Table 4.1: State-Space Form of Zone Lumped Model

zone tracking errors, and the total amount of cooling delivered. The inputs are the zone temperature setpoints. The disturbances are the ambient temperature, external loads, and steady-state cooling rates.

The energy balance and heat duty model can be expressed as a continuous-time state-space model with $D = 0$.

$$\frac{dx}{dt} = A_c x + B_c u + \Gamma_c d$$
$$y = Cx + Du$$

The continuous-time model above is discretized using the controller interval, Δ, to get the discrete version of the matrices: A, B, and B_d. The final discrete-time state space model used in SI is given by

$$x^+ = Ax + Bu + B_d d$$
$$y = Cx + Du$$

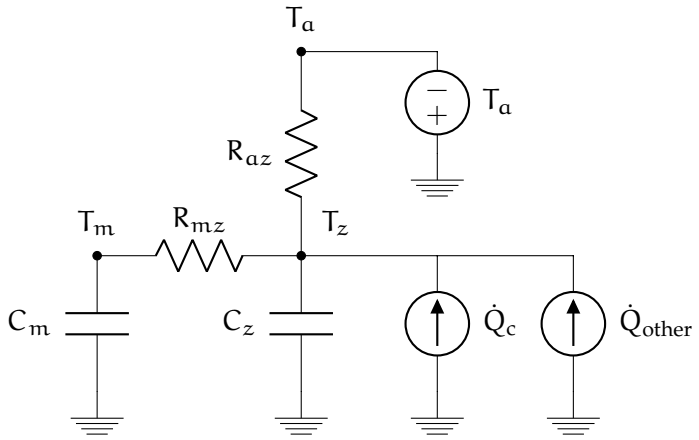

Figure 4.4: Network for zone air/mass model where $R_{(\cdot)} = \frac{1}{H_{(\cdot)}}$.

4.4.2 Single Zone Air/Mass Model

For the single zone air/mass model, separate states are used to model the zone air and mass temperatures. With a PI controller, the model is given by

$$\begin{aligned}
C_{z,i}\frac{dT_{z,i}}{dt} &= H_{az,i}(T_a - T_{z,i}) + H_{mz,i}(T_{m,i} - T_{z,i}) - \dot{Q}_{c,i} + \dot{Q}_{other,i} \\
C_{m,i}\frac{dT_{m,i}}{dt} &= H_{mz,i}(T_{z,i} - T_{m,i}) \\
\dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right] \\
\varepsilon_i &= T_{sp,i} - T_{z,i}
\end{aligned}$$

A depiction of the corresponding thermal network is shown in Figure 4.4, however the heat transfer form is used in the document.

As shown in Table 4.2, the states are the zone air and mass temperatures, the integrals of the zone tracking errors, and the total amount of cooling delivered. The inputs are the zone temperature setpoints. The disturbances are the ambient temperature, external loads, and steady-state cooling rates. A similar procedure as before can be followed to obtain a discrete-time state space model.

States	$x = \begin{bmatrix} T_{z,i} & T_{m,i} & \bar{\varepsilon}_i & Q_{total} \end{bmatrix}^T$	$n = 3n_{bz} + 1$
Inputs	$u = \begin{bmatrix} T_{sp,i} \end{bmatrix}^T$	$m = n_{bz}$
Measurements	$y = \begin{bmatrix} T_{z,i} & Q_{total} \end{bmatrix}^T$	$p = n_{bz} + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,i} & \dot{Q}_{ss,i} \end{bmatrix}^T$	$n_d = 1 + 2n_{bz}$

Table 4.2: State-Space Form of Zone Air/Mass Model

States	$x = \begin{bmatrix} T_{z,i} & T_{s,i} & T_{d,i} & \bar{\varepsilon}_i & Q_{total} \end{bmatrix}^T$	$n = 4n_{bz} + 1$
Inputs	$u = \begin{bmatrix} T_{sp,i} \end{bmatrix}^T$	$m = n_{bz}$
Measurements	$y = \begin{bmatrix} T_{z,i} & Q_{total} \end{bmatrix}^T$	$p = n_{bz} + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,i} & \dot{Q}_{ss,i} \end{bmatrix}^T$	$n_d = 1 + 2n_{bz}$

Table 4.3: State-Space Form of Zone Air/Shallow/Deep Model

4.4.3 Single Zone Air/Shallow/Deep Model

For the single zone air/mass model, separate states are used to model the zone air, shallow mass, and deep mass temperatures. With a PI controller, the model is given by

$$
\begin{aligned}
C_{z,i}\frac{dT_{z,i}}{dt} &= H_{az,i}(T_a - T_{z,i}) + H_{sz,i}(T_{s,i} - T_{z,i}) - \dot{Q}_{c,i} + \dot{Q}_{other,i} \\
C_{s,i}\frac{dT_{s,i}}{dt} &= H_{sz,i}(T_{z,i} - T_{s,i}) + H_{ds,i}(T_{d,i} - T_{s,i}) \\
C_{d,i}\frac{dT_{d,i}}{dt} &= H_{ds,i}(T_{s,i} - T_{d,i}) \\
\dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right] \\
\varepsilon_i &= T_{sp,i} - T_{z,i}
\end{aligned}
$$

Likewise, a depiction of the corresponding thermal network is shown in Figure 4.5 for comparison with other texts in this field.

As shown in Table 4.3, the states are the zone air, shallow mass, and deep mass

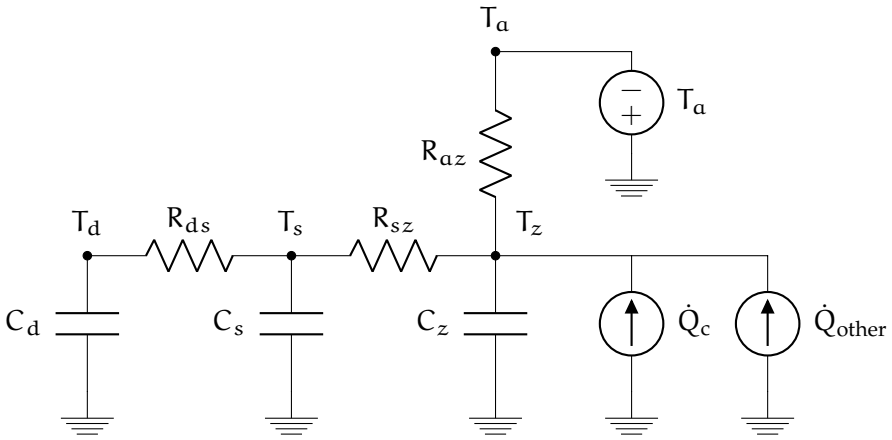

Figure 4.5: Network for zone air/shallow/deep model where $R_{(\cdot)} = \frac{1}{H_{(\cdot)}}$.

temperatures; the integrals of the zone tracking errors; and the total amount of cooling delivered. The inputs are the zone temperature setpoints. The disturbances are the ambient temperature, external loads, and steady-state cooling rates.

Coupling Between Zones

For multiple zones, coupling is modeled between the zone (air) temperatures.

$$C_{z,i}\frac{dT_{z,i}}{dt} = H_{az,i}(T_a - T_{z,i}) + H_{sz,i}(T_{s,i} - T_{z,i}) + \sum_{j\neq i}\beta_{ij}(T_{z,j} - T_{z,i}) - \dot{Q}_{c,i} + \dot{Q}_{other,i}$$

Depending on the physical properties of the building, the coupling can also be modeled between shallow mass temperatures.

$$C_{s,i}\frac{dT_{s,i}}{dt} = H_{sz,i}(T_{z,i} - T_{s,i}) + H_{ds,i}(T_{d,i} - T_{s,i}) + \sum_{j\neq i}\beta_{ij}(T_{s,j} - T_{s,i})$$

Table 4.4: SI Model Parameters

States	$x = \begin{bmatrix} T_{z,i} & \bar{\varepsilon}_i & Q_{total} \end{bmatrix}^T$	$n = 2n_{bz} + 1$
Inputs	$u = \begin{bmatrix} T_{sp,i} \end{bmatrix}^T$	$m = n_{bz}$
Measurements	$y = \begin{bmatrix} T_{z,i} & Q_{total} \end{bmatrix}^T$	$p = n_{bz} + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,i} & \dot{Q}_{ss,i} \end{bmatrix}^T$	$n_d = 1 + 2n_{bz}$
Physical Parameters	$\{k_i, (mC_p)_i, \beta_{ij}\}$	$2n_{bz} + \begin{pmatrix} n_{bz} \\ 2 \end{pmatrix}$
Controller Parameters	$\{K_{c,i}, \tau_{I,i}\}$	$2n_{bz}$

4.5 SI Approach

For SI, the lumped model is used as the grey-box model form. The model variables and list of parameters to be identified for the grey-box model are given in Table 4.4. The black-box model does not assume any specific structure, so it must compute the entire state-space model matrices A, B, and C.

In addition to input to output relationships, the effects of measurable disturbances can also be identified during the SI phase. One of the primary disturbances affecting buildings, ambient temperature, is easily measured and predicted. Weather forecasts can be used during online optimization if the heat transfer coefficient with the ambient is computed by adding ambient temperature to the input vector before passing the data to the identification routine. The model between other disturbances, such as solar radiation, and the zone temperature can also be identified if measurements and predictions of these disturbances are available.

4.6 SI Results

To determine how well the black-box and grey-box identification perform, both methods are applied to the same dataset using MATLAB's System Identification Toolbox.

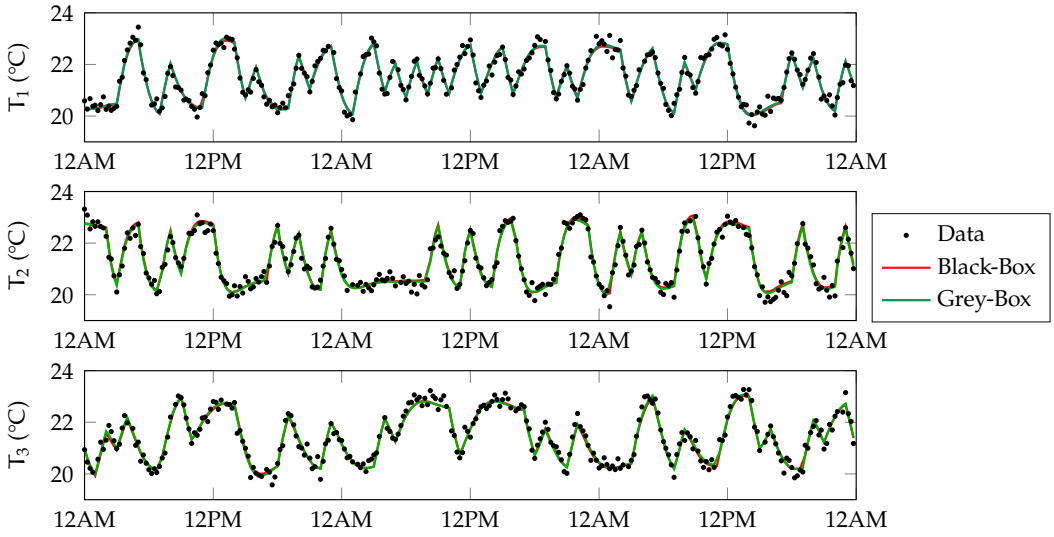

Figure 4.6: Validation fit from system identification

The free parameters in the grey-box approach include the thermal capacitance of the zones and heat transfer coefficients as well as the PI controller parameters The results from this study are shown in Figure 4.6. The identified models are used to generate predictions for validation data not used for training. Based on the fits, both methods capture the behavior the validation data well. They appear nearly on top of one another. Hence, both approaches should be suitable for use in applications. However, the two methods do not scale equivalently for large-scale systems as examined next.

4.7 Scaling of Free Parameters

As the number of zones increases, the number of free parameters that must be identified using SI also increases. Figure 4.7 shows this scaling for both grey-box and black-box models. The lumped model is selected for the grey-box category, and the plot depicts the *worst case* scenario for grey-box identification when every heat

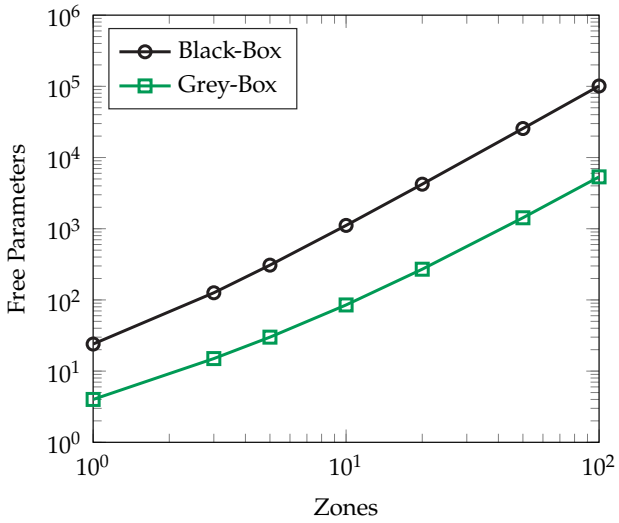

Figure 4.7: Scaling of free parameters for grey-box and black-box identification

transfer coefficient between each pair of zones need to be determined. Typically, many of these coupling coefficients β_{ij} can be set to zero a priori when zones are not adjacent. Doing so leads to a more linear scaling between free parameters and the number of zones for the grey-box case.

From this figure, it is evident that black-box identification has many more free parameters that must be determined. Since large-scale commercial building applications may have hundreds of zones, developing black-box models for these systems involves identifying many parameters. In fact, more identification data is required when there are more free parameters to identify. Thus, grey-box identification requires significantly less data than black-box identification. In practice, less data means a shorter testing period and faster implementation of the MPC controller, which makes grey-box identification more suitable for large-scale commercial applications. Instead of testing periods of several weeks, identification can be performed with just a couple days worth of data which can provide higher rates of return on investment. These findings were verified via simulation studies by decreasing the

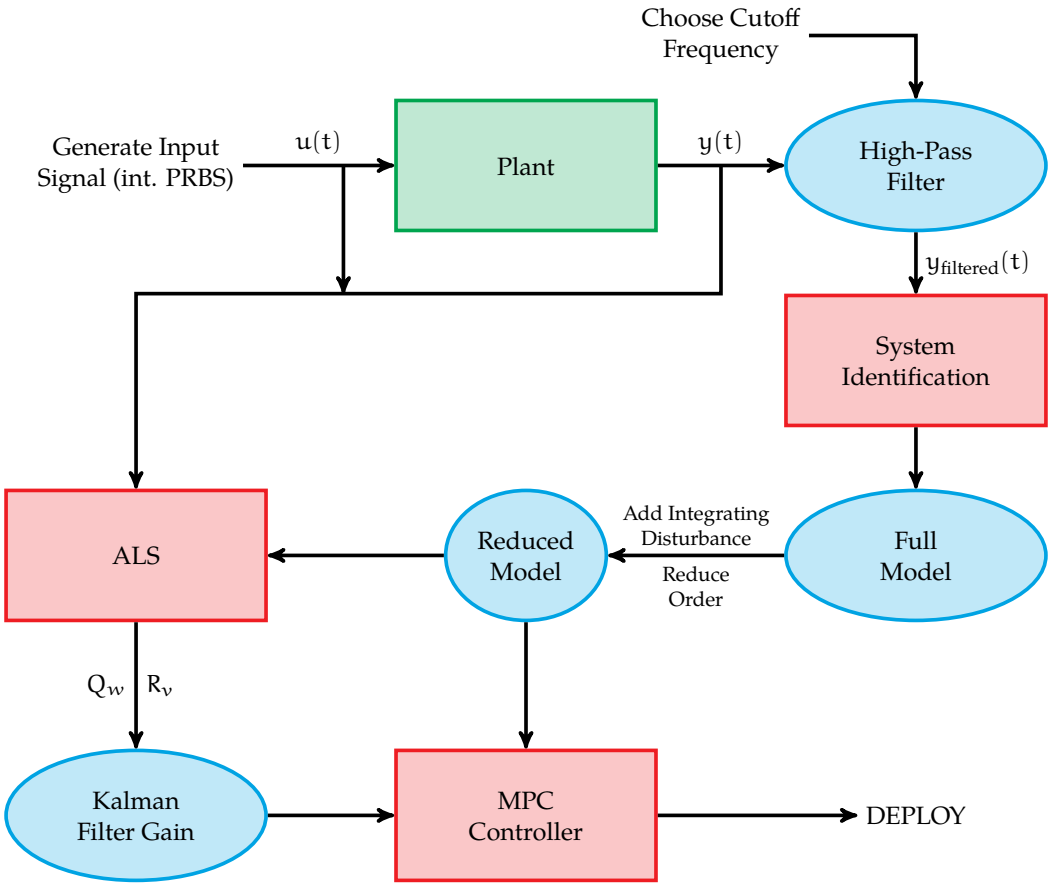

Figure 4.8: Model development flow fiagram

size of the training set. The grey-box approach performed well with little data while the black-box approach would result in poor validation fits.

4.8 Deployment Strategy

In practice, there are several steps that are needed to deploy the MPC architecture. Some of these steps are highlighted in Figure 4.8. Identification tests are run by exciting the inputs (temperature setpoints) and measuring the outputs (temperatures and cooling duties). The output data can be filtered to remove the slow frequency disturbances if required by selecting a cutoff frequency and using a high-pass filter.

This optional step may be needed when there are drifts in the data caused by seasonal variations. This filtered data can then be used for SI, as discussed with grey-box techniques. The model can be reduced if the order is too high and there are extra states that are not necessary for capturing the main dynamics of the system. A lumped model would have minimal states, but perhaps an air/shallow/deep mass representation can be simplified further.

With the reduced model, integrating disturbance states are added to achieve zero offset performance in the present of plant model mismatch and unmeasured disturbances [83, Ch. 1]. The reduced model along with the test data can be sent to an autocovariance least-squares (ALS) method to identify the noise covariance matrices [63, 98]. These noise covariances are used to obtain the Kalman Filter gain for the state estimator. With the identified model and Kalman Filter gain, the MPC controller is then ready to be deployed.

4.9 Airside Power Consumption Models

Until now, addressing power consumption by waterside equipment has been the primary concern. However, airside equipment, such as fans in AHUs, also consume power. Historically, these contributions have been small compared to power consumption at central plants and have been ignored. In this section, the magnitude of these contributions are explored to determine if they are significant enough to include in the control architecture.

4.9.1 Efficiency Calculations

Approximate calculation of the efficiency is first performed using typical values for the airside system.

Cooling Capacity of Air For air at 20°C, the density and heat capacity can be obtained from a thermodynamics handbook as

$$\rho_{air} = 1.205 \ \frac{kg}{m^3} \qquad C_{p,air} = 1.005 \ \frac{kJ}{kg \cdot K}$$

Assuming that the supply air from the AHU is at 55°F and the room air is at 72°F, the cooling capacity of air is computed using

$$\begin{aligned}
\dot{Q}_c &= \dot{m}_{air} C_{p,air} \Delta T = \rho_{air} \dot{V}_{air} C_{p,air} (T_{room} - T_{supply}) \\
\frac{\dot{Q}_c}{\dot{V}_{air}} &= \rho_{air} C_{p,air} (T_{room} - T_{supply}) \\
\frac{\dot{Q}_c}{\dot{V}_{air}} &= \left(1.205 \ \frac{kg}{m^3}\right) \left(1.005 \ \frac{kJ}{kg \cdot K}\right) (72°F - 55°F) \left(\frac{5 \ K}{9 \ °F}\right) \\
\frac{\dot{Q}_c}{\dot{V}_{air}} &= 11.437 \ \frac{kJ}{m^3}
\end{aligned}$$

Power Consumption Rules of Thumb To get a rough estimate of power consumption, typical HVAC rules of thumb can be employed [1]. In particular, the two values shown below are used

Fan Energy	1000–1500 CFM/hp
Chillers & Pumps & Towers	0.9–1.0 kW/ton

Hence to provide 1 ton of refrigeration, 0.9–1.0 kW of electricity for the waterside operation (chillers, pumps, and towers) is required. To provide the same 1 ton of refrigeration, power must also be supplied to fans in the air handler units. This power can be estimated using the cooling capacity of air and the rule of thumb for

	Power Requirement (kW/ton)	Percentage of Total Cost
Airside	0.32–0.49	24–35%
Waterside	0.9–1.0	65–76%

Table 4.5: Summary of Power Consumption Calculations

fan energy as

$$\begin{aligned}
1 \text{ ton} &= 3.517 \text{ kW cooling} \left(\frac{\dot{V}_{\text{air}}}{\dot{Q}_c} \right) = 3.517 \text{ kW cooling} \left(\frac{1 \text{ m}^3}{11.437 \text{ kJ}} \right) \\
&= 0.3075 \frac{\text{m}^3}{\text{s}} \left(\frac{2118.9 \text{ CFM}}{1 \text{ m}^3/\text{s}} \right) \left(\frac{1 \text{ hp}}{1000\text{–}1500 \text{ CFM}} \right) \\
&= 0.4344\text{–}0.6515 \text{ hp} \left(\frac{0.7457 \text{ kW}}{1 \text{ hp}} \right) \\
&= 0.32\text{–}0.49 \text{ kW}
\end{aligned}$$

Hence, 0.32–0.49 kW is required for operating fans. The results are summarized in Table 4.5. As a rough estimate, airside operation accounts for 30% of the total electricity cost and waterside accounts for 70%. Based on these results, clearly airside power consumption must be considered in the economic optimization since it comprises a significant portion of the total electricity cost.

Efficiency Estimates Based on the results shown in Table 4.5, estimates of the coefficient of performance, a measure of efficiency as defined in Equation (2.3), can be obtained for both waterside and airside operation

$$\begin{aligned}
\text{COP}_{\text{waterside}} &= \left(\frac{3.517}{0.9\text{–}1.0 \text{ kW/ton}} \right) &&= 3.5\text{–}3.9 \\
\text{COP}_{\text{airside}} &= \left(\frac{3.517}{0.32\text{–}0.49 \text{ kW/ton}} \right) &&= 7.2\text{–}11
\end{aligned}$$

$$\eta_{HVAC} = \frac{1}{COP_{waterside}} = \left(\frac{1}{3.5\text{–}3.9}\right) \quad = 0.26\text{–}0.28$$
$$\eta_{air} = \frac{1}{COP_{airside}} = \left(\frac{1}{7.2\text{–}11}\right) \quad = 0.091\text{–}0.14$$

One set of values used in simulations are $\eta_{HVAC} = 0.27$ and $\eta_{air} = 0.1$.

As a first step towards developing an airside power consumption model, the relationship between cooling load and airside power requirement is assumed to be linear; therefore, the efficiency, η_{air}, is a constant. If this approximation is not sufficiently accurate, η_{air} can be determined as a function of load and other parameters making the high-level problem a nonlinear optimization. Industrial large-scale aggregate models from practice are required to validate this assumption.

4.9.2 Modification of High-Level Problem

With the inclusion of the airside power consumption model, the high-level MPC problem is reformulated as

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{u},\dot{W}_{peak}} \quad & \sum_{k=0}^{N-1} c_k \left(\eta_{HVAC}\dot{Q}_{HVAC,k} + \eta_{air}\sum_b \dot{Q}_{c,b,k}\right)\Delta + c_{peak}\dot{W}_{peak} \\
\text{s.t.} \quad & C_b\frac{dT_b}{dt} = -H_b(T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{other,b} \\
& \frac{ds}{dt} = -\sigma s + \dot{Q}_{storage} \\
& \dot{Q}_{HVAC,k} = \sum_b \dot{Q}_{c,b,k} + \dot{Q}_{storage,k} \\
& \left(\eta_{HVAC}\dot{Q}_{HVAC,k} + \eta_{air}\sum_b \dot{Q}_{c,b,k}\right) \leqslant \dot{W}_{peak} \\
& \dot{W}_{peak,past} \leqslant \dot{W}_{peak}, 0 \leqslant \dot{Q}_{HVAC,k} \leqslant \dot{Q}_{HVAC,max} \\
& 0 \leqslant s_k \leqslant s_{max} \\
& T_{min} \leqslant T_b \leqslant T_{max}
\end{aligned}$$

The airside operation determines η_{air} and the waterside operation determines η_{HVAC}, both of which are treated as constants in the High Level problem and can be updated periodically via feedback based on the operation of airside and waterside equipment. The demand change constraint is also changed to reflect this contribution from both subsystems. Note that active TES is not able to shift this component of the power load. Even if chillers are run overnight charging the active TES tank, power is still required during the peak periods to run the airside equipment in order to deliver the cooling load to zones.

4.10 Summary

The development of models is a key step in MPC deployment. For commercial HVAC systems, a grey-box system identification procedure is recommended over the black-box approach to reduce the amount of training data required and subsequently costs of commissioning controllers. With more efficient central plant equipment technology available, airside power consumption is becoming a larger fraction of the total electricity bill, so it can no longer be ignored. Since active TES cannot shift this portion of the power load, other energy storage are necessary. One way to shift this airside power load is to use electrical storage, which is explored in the next chapter.

5 EXTENSIONS

Thermal energy storage is able to shift the cooling load but not necessarily the power load. Although production of chilled water is the primary source of power consumption in central HVAC plants, airside equipment, such as fans in air handler units, also consume power when delivering the produced chilled water to zone units for cooling. In order to shift this electrical load, electrical storage is needed. In this chapter, we extend the framework to consider batteries as a method for load shifting in HVAC applications. VRF systems are also discussed. These extensions are used to demonstrate the wide applicability of the proposed decomposition strategy.

5.1 Batteries

When discussing batteries, there are many technologies available [99]. With electrical storage, the power purchased during off-peak periods is stored in batteries which are discharged when power is needed by HVAC equipment [6, 49]. While TES is often the more commonly considered storage method for load shifting applications due to its high efficiency and low capital costs, electrical storage allows a larger portion of the load to be shifted than TES.

In the past, the cost of manufacturing batteries has been high, however, prices have been decreasing over the past decade [62]. Additionally, batteries have the added benefit that they can be used for revenue generation on the electricity market. Even when not load shifting, they can be used to generate income by buying electricity when cheap and selling it back to the market when expensive [19]. In this section, the large-scale embedded battery problem is addressed.

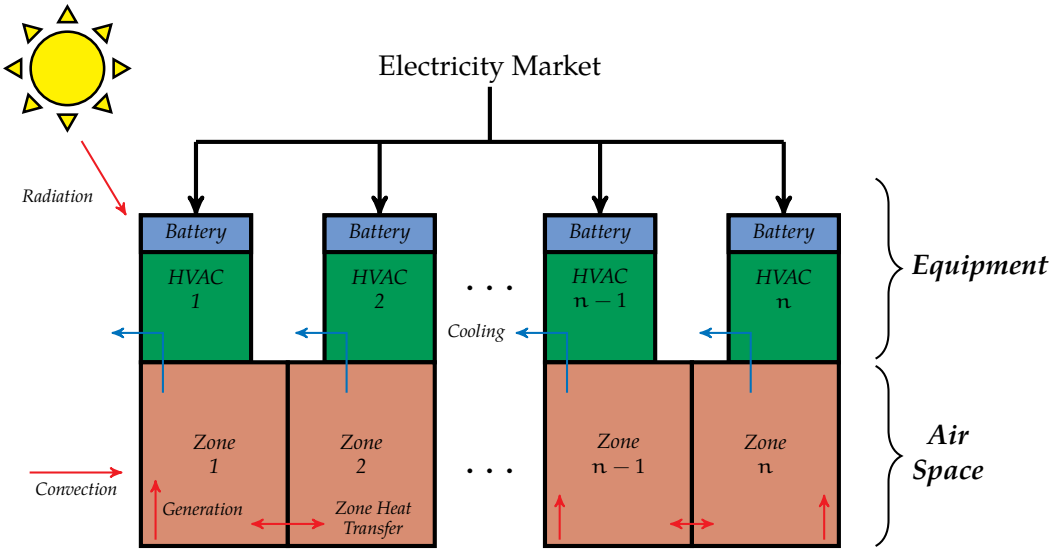

Figure 5.1: Embedded battery HVAC applications

5.1.1 Distributed Embedded Battery Units

The embedded battery system considered in this application is depicted in Figure 5.1: there is a collection of n zones, each regulated using an HVAC unit. Zones may or may not exchange heat with each other depending on if they share a wall [67]. Each of these zone HVAC units has an embedded battery associated with it as well as a thermostat than can receive zone setpoints. The HVAC units can draw power from the battery or the electricity market. The control objective is to minimize the total energy cost. The decisions made by the control system include the zone temperature setpoints and the charging or discharging schedule of the embedded batteries. Embedded batteries are not shared between different units. Constraints on equipment capacities and occupancy comfort must be respected.

5.1.2 Decomposition

To address this class of HVAC systems, the proposed hierarchical decomposition is shown in Figure 5.2 [71]. Notice that the structure is similar to the decomposition

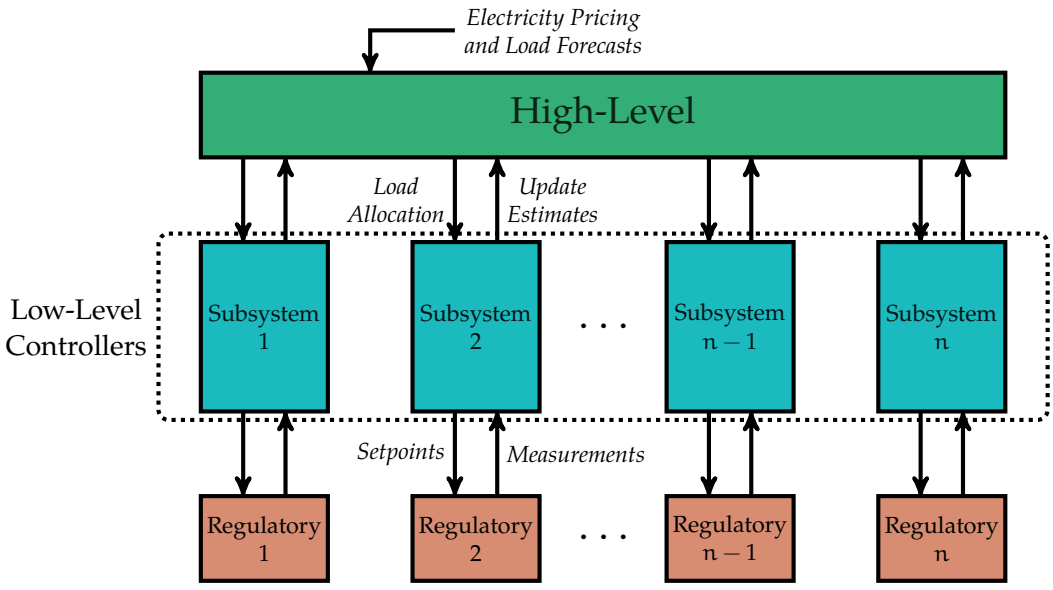

Figure 5.2: Hierarchical decomposition for distributed embedded battery units

presented in Figure 3.2 without the waterside problem since there is no central HVAC plant in this application. As before, the collection of zones is grouped into a series of subsystems. Division by building is the choice we make to decouple the subsystems. Corresponding models are created to describe the corresponding aggregate subsystem dynamics.

Load and pricing forecasts serve as inputs to the high-level controller, which performs a system-wide economic optimization using these aggregate subsystem models to reduce the number of variables in the optimization problem compared to the combined problem with full detailed zone models. The high-level controller determines the allocation of power to each low-level subsystem. In each subsystem, the low-level problem is solved to obtain the control policy that minimizes energy usage subject to not exceeding the load allocated from the high-level problem using detailed zone models. These low-level problems can be solved in parallel since they depend only on the high-level problem for the allocation and do not need to

communicate with other low-level controllers.

Mathematical formulations for these problems are discussed in Sections 5.1.4 and 5.1.5. To ensure a feasible policy, necessary constraints are added to the optimizations. The computed decisions are sent from the low-level controllers to the regulatory layer (zone thermostats) in which they are implemented. At the next controller execution, measurements are collected from the regulatory layer and used as feedback to correct for model mismatch and unmeasured disturbances as discussed in previous chapters.

Benefits of this Structure Decomposition using this structure has several benefits, similar to the ones for the decomposition in Figure 3.2. The peak demand charge, which links all subsystems through the economic objective, is addressed in the high-level problem where the system-wide economic optimization is performed. The computational load is reduced via the use of aggregate models in the high-level problem. Without required iterations between subsystem controllers, the communication requirements are minimal; the low-level problems can be solved in parallel. As a result, the architecture can be implemented in real time. The architecture is easily scalable to handle large embedded battery applications with many zones.

5.1.3 Battery Model

The lumped zone model from Chapter 4 is used for describing temperature dynamics. Dynamics for electrical storage evolution is also needed for the MPC controllers. A first-order linear model with decay is used for batteries. The zone embedded battery

level, B_i, evolves according to

$$\frac{dB_i}{dt} = -\sigma_i B_i + \sqrt{\eta_{batt,i}}\dot{W}_{in,i} - \frac{1}{\sqrt{\eta_{batt,i}}}\dot{W}_{out,i} \tag{5.1}$$

in which σ represents the decay term, $\eta_{batt,i}$ represents the round-trip efficiency, and $\dot{W}_{in,i}, \dot{W}_{out,i}$ represent the charging and discharging rates. Analogously, an aggregate electrical storage model can be expressed at the subsystem level using the subscript b to denote buildings in place of i for zones.

5.1.4 High-Level Problem Formulation

With the models defined, high-level problem is formulated as an optimization problem as given in Equation (5.2). With a horizon of N steps, the economic objective consists of two components: time-dependent usage prices c_k and the peak demand charge c_{peak}. The efficiency of the HVAC units, η_{HVAC}, is used as an aggregate measure of equipment performance. Zone battery models are consolidated into aggregate subsystem electrical storage models. Several constraints are included for power demand satisfaction, peak demand charge, HVAC unit capacities, comfort bounds on temperatures and battery capacities. The comfort constraints are enforced as soft constraints by penalizing their violations in the objective to maintain feasibility in the presence of large disturbances that cannot be reject by the equipment. The power allocation computed from this problem is sent to each of the low-level problems.

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u},\dot{W}_{\text{peak}}} \quad & \sum_{k=0}^{N-1} c_k \dot{W}_{\text{grid},k} \Delta + c_{\text{peak}} \dot{W}_{\text{peak}} \\
\text{s.t.} \quad & C_b \frac{dT_b}{dt} = -H_b (T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{\text{other},b} \\
& \frac{dB_b}{dt} = -\sigma B_b + \sqrt{\eta_{\text{batt},b}} \dot{W}_{\text{in}} - \frac{1}{\sqrt{\eta_{\text{batt},b}}} \dot{W}_{\text{out}} \\
& \dot{W}_{\text{HVAC},k} = \eta_{\text{HVAC}} \sum_b \dot{Q}_{c,b,k} \\
& \dot{W}_{\text{grid},k} = \dot{W}_{\text{HVAC},k} + \dot{W}_{\text{in},k} - \dot{W}_{\text{out},k} \leqslant \dot{W}_{\text{peak}} \\
& \dot{W}_{\text{peak,past}} \leqslant \dot{W}_{\text{peak}}, 0 \leqslant \sum_b \dot{Q}_{c,b,k} \leqslant \dot{Q}_{\text{HVAC,max}} \\
& 0 \leqslant B_{k,b} \leqslant B_{\text{max},b} \\
& T_{\text{min}} \leqslant T_b \leqslant T_{\text{max}}
\end{aligned}
\tag{5.2}
$$

5.1.5 Low-Level Problem Formulation

The low-level optimization problem is formulated for each subsystem as shown in Equation (5.3). The objective is to minimize the total energy usage while not exceeding the allocation from the high-level problem $\dot{W}_{\text{HighLevel},k}$. The detailed zone models are used in this problem and analogous constraints are enforced for comfort and capacity. The computed setpoints from this problem are sent to the zone thermostats.

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u}} \quad & W_{\text{total},N} \\
\text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \\
& \frac{dB_i}{dt} = -\sigma_i B_i + \sqrt{\eta_{\text{batt},i}}\dot{W}_{\text{in},i} - \frac{1}{\sqrt{\eta_{\text{batt},i}}}\dot{W}_{\text{out},i} \\
& \dot{Q}_{c,i} = f(T_i, T_{\text{sp},i}) \\
& \frac{dW_{\text{total}}}{dt} = \sum_i \eta_{\text{HVAC},i}\dot{Q}_{c,i} + \dot{W}_{\text{in},i} - \dot{W}_{\text{out},i} \\
& T_{\min} \leqslant T_i \leqslant T_{\max} \\
& 0 \leqslant B_{k,i} \leqslant B_{\max,i} \\
& \frac{W_{\text{total},k+1} - W_{\text{total},k}}{\Delta} \leqslant \dot{W}_{\text{HighLevel},k} \\
& W_{\text{total},k+1} - W_{\text{total},k} \geqslant 0
\end{aligned}
\tag{5.3}
$$

5.1.6 Simulation Results

The proposed control decomposition is applied to a campus of 12 buildings, with each building containing 8 zones. Each zone HVAC unit in this 96-zone system has its own separate embedded battery. The total electrical storage available is 60 MWh, spread evenly over all zone units. Ambient temperature and electricity pricing data used for the simulation are displated in Figure 5.3.

The optimization problems were solved using Gurobi 7.5.2 via MATLAB R2017b on a desktop machine with 16GB RAM and 2.70GHz Intel® Core™ Processor i5-7500T. The high-level problem was solved in 1.3 seconds, while each of the 12 low-level problems was solved in approximately 1.6 seconds. Since these low-level problems can be solved in parallel, the required computational time is 3 seconds. Hence, this architecture can be implemented online for applications for which the MPC time step is typically 15 min.

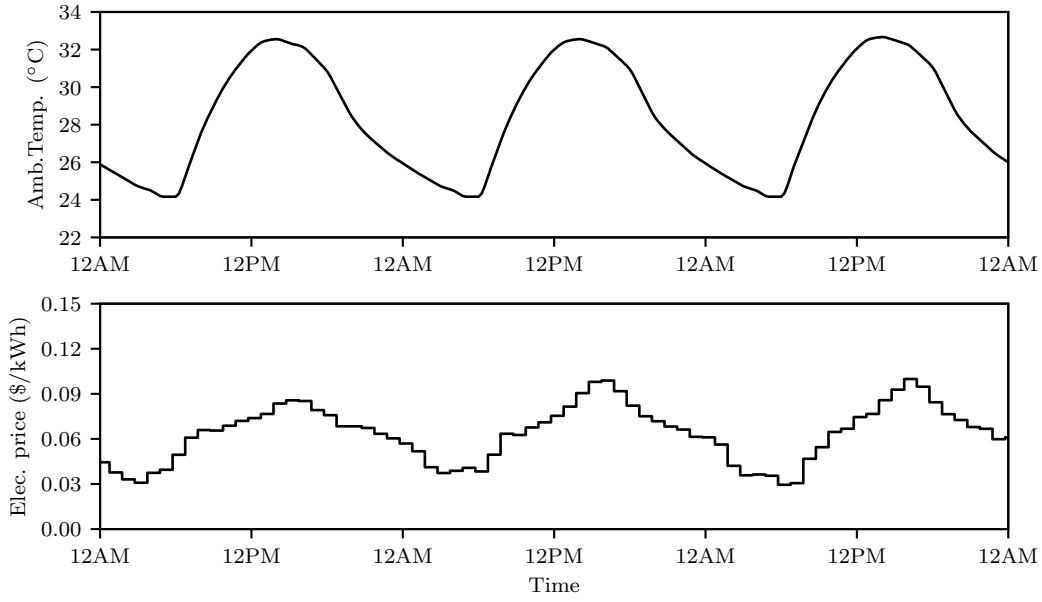

Figure 5.3: Weather and pricing data used in the embedded batteries simulation

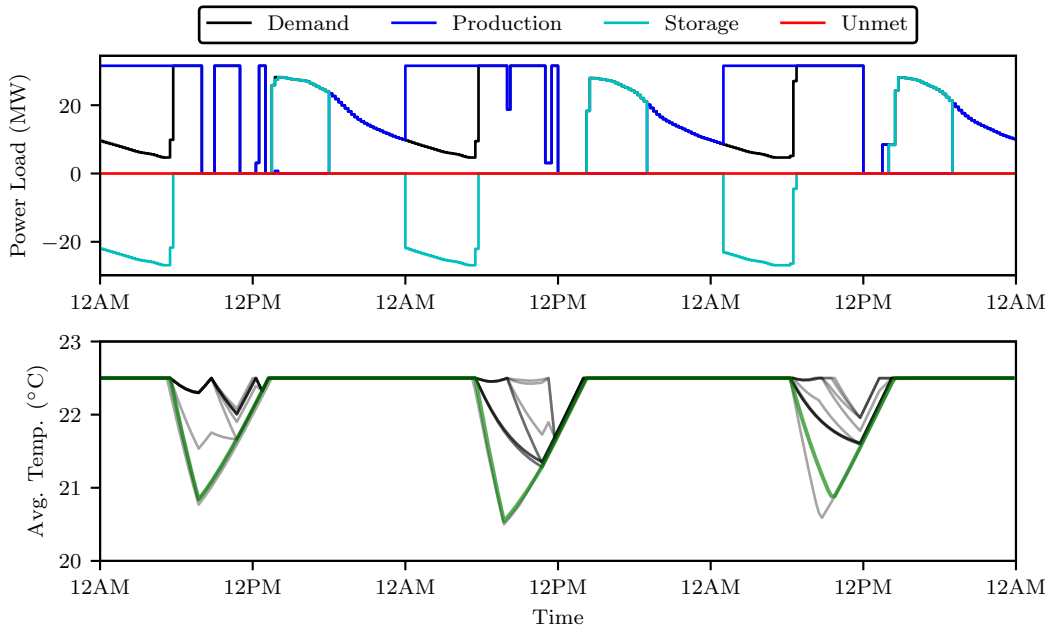

Figure 5.4: Results from the high-level embedded batteries optimization

The high-level problem results are plotted in Figure 5.4. The top plot shows the total HVAC power consumption across the campus for the three-day period and how this load is met using a combination of battery storage and power purchase from the market. The second plot shows the 12 average building temperature profiles, with one particular building highlighted in green. The optimizer uses both batteries and building mass for load shifting. For electrical storage, the batteries are charged overnight and then discharged in the late afternoon so less power is purchased from the market during the peak periods to meet the demand. With passive TES, the buildings are precooled in the early morning so less cooling is needed in the middle of the day. The peak demand charge is also optimized in the high-level problem to prevent the load shifting from leading to a large spike in power usage, resulting in the flatness of the total production plot. These power load allocations are sent to the low-level problems.

Figure 5.5 plots the results computed for the 96 zones from the 12 low-level problems. The zone temperatures, setpoints, and embedded battery usage profiles are shown in black with one particular zone highlighted in green. The aggregate energy storage strategy outlined by the high-level is followed with added resolution due to the detailed zone models. The time step is 15 min. However, if rapid changes between charging and discharging are not desired on the order of 15 min, additional rate-of-change constraints and penalties can be added to obtain the desirable behavior. The zone temperature setpoints are sent to the regulatory layer.

The total operational cost for this three-day period is $37.5K using the proposed approach. For comparison with a baseline, the cost of the energy minimization solution (staying at the upper bound of the comfort region without load shifting) is $45.4K, hence load shifting is able to provide about 18% savings using the passive TES and battery storage. When the amount of available electrical storage is increased,

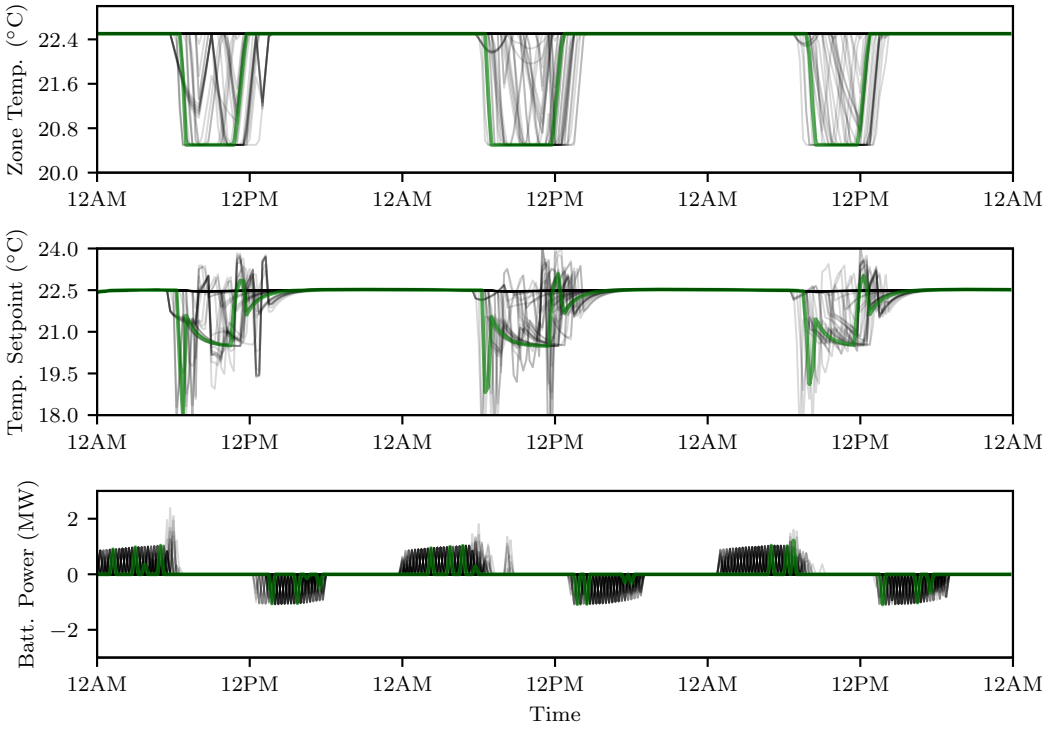

Figure 5.5: Results from the low-level embedded batteries optimization

these cost savings can be as high as 25% in the particular system and pricing structure considered here. With additional economic incentives, the cost savings number can rise to 40%.

5.1.7 Tradeoff between Cost and Comfort

A study was performed to analyze the tradeoff between occupancy comfort and energy costs. For this study, the temperature setpoint was varied from the upper bound to the middle of the comfort zone in the baseline control strategy. The energy cost was computed for each point in the analysis without load shifting. The resulting tradeoff curve between cost and comfort is given in Figure 5.6. The total energy cost is plotted against the average squared temperature deviation away from 21.5 °C. As

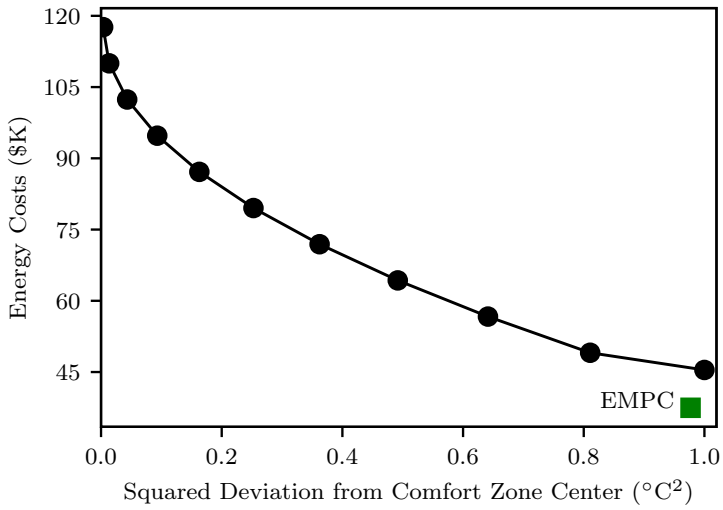

Figure 5.6: Tradeoff curve between energy cost and occupancy comfort

comfort increases, costs rise. The price of comfort is not cheap for large systems due to the time-varying usage electricity prices as well as the peak demand charge. The green square denotes the solution of the hierarchical decomposition scheme presented earlier. It achieves a lower cost as well as a smaller comfort violation compared to the energy minimization solution of staying at the upper bound of the comfort region. This result suggests that load shifting reduces energy costs without increasing discomfort.

5.1.8 Economics of Battery Storage

Though electrical storage can provide additional savings, capital investment is needed to purchase batteries. In this section, an economic analysis is presented of the usage of batteries for load shifting in HVAC systems in order to approximate the financial feasibility of this approach. This economic analysis is also presented in [73]. In order for embedded battery units to compete with conventional HVAC units, the added savings must offset the extra capital cost required to provide electrical

storage.

5.1.9 Assumptions

Two primary assumptions are made in this estimate.

First, for load shifting applications, the savings are achieved by buying power at lower off-peak prices so less power is required to be purchased at on-peak prices. For the electricity pricing data used in the study, the difference in these prices is on the order of $0.05 per kWh. However, this margin may be higher in other electricity markets. Batteries also provide additional savings through demand response and managing the peak demand charge. However, there may be conversion losses due to the round-trip efficiency being less than 100% in electrical storage which reduce savings. A (generous) cost savings of *$0.10/kWh* of battery storage is assumed.

Second, the energy cost savings also depend on the number of times the battery is charged and discharged. One cycle per day (*365 cycles per year*) is assumed since one day is the natural period in load shifting applications; however, there may be fewer cycles if the building does not operate 365 days per year. More cycles cannot be achieved since the cost savings rely on charging during the off-peak period and discharging during the peak period which occur only daily.

5.1.10 Large High Performance Batteries

We begin by considering large-scale batteries. Capital pricing data for two batteries obtained from Johnson Controls is given in Table 5.1. C-rate is a measure of how quickly the battery discharges its nominal capacity. It can be seen as the ratio of power to capacity in units of inverse hours. A 1C rate battery discharges completely in one hour at nominal power output; a 2C rate battery discharges completely in 30 min.

Table 5.1: Capital Investment for High Performance Batteries

	Battery A	**Battery B**
Capacity	2000 kWh	500 kWh
Power	2000 kW	1000 kW
C-rate	1	2
Cost	\$1256K	\$763K

Based on the assumptions made, the annual cost savings for these two batteries are

$$\text{Battery A:} \qquad \left(\frac{\$\ 0.10}{1\ \text{kWh}}\right)\left(\frac{2000\ \text{kWh}}{1\ \text{cycle}}\right)\left(\frac{365\ \text{cycles}}{1\ \text{year}}\right) = \$73{,}000\ \text{per year}$$

$$\text{Battery B:} \qquad \left(\frac{\$\ 0.10}{1\ \text{kWh}}\right)\left(\frac{500\ \text{kWh}}{1\ \text{cycle}}\right)\left(\frac{365\ \text{cycles}}{1\ \text{year}}\right) = \$18{,}250\ \text{per year}$$

Since Battery A has a capital investment of \$1256K with annual savings of \$73K, its payback period is approximately *17 years*. Similar for Battery B with a capital investment of \$763K and annual savings of \$18.25K, its payback period is approximately *42 years*. Neither would be financially attractive since the batteries may have a lifetime of only 10 years.

5.1.11 Electric Car Batteries

To make these payback periods more attractive, the capital investment costs must be decreased. A large portion of the capital cost of Batteries A and B is attributed to the high power requirement of these two batteries (high C-rates). For load shifting, it may not be necessary to discharge the battery completely in 30 minutes or 1 hour. It is sufficient to use batteries with a lower C-rate that can discharge in 2–4 hours (0.25C–0.5C) since peak and off-peak periods are of this order of magnitude.

As an example, commercial electric car batteries are explored. An assembly of

Table 5.2: Economic Analysis Results

	Battery A	**Battery B**	**Car Battery**
Cost/kWh	\$628	\$1526	\$145
Savings	\$73K	\$18.25K	\$36.50
PP	17 years	42 years	4 years

these batteries can be used in parallel to meet the electrical storage requirements of commercial HVAC systems [34]. Additionally, applications with distributed embedded batteries may have many of these smaller batteries instead of a small number of large ones, which can help in reducing their footprint.

In 2016, the battery cell price for the Chevrolet Bolt was \$145 per kWh [24]. This cost is projected to decrease to below \$100/kWh over the next decade [8]. Using a price of \$145/kWh, the savings for each kWh battery cell are

$$\text{Car Battery Cell:} \qquad \left(\frac{\$\ 0.10}{1\ \text{kWh}}\right)\left(\frac{1\ \text{kWh}}{1\ \text{cycle}}\right)\left(\frac{365\ \text{cycles}}{1\ \text{year}}\right) = \$36.50\ \text{per year}$$

Since each kWh car battery cell has a capital investment of \$145 with annual savings of \$36.50, its payback period is approximately *4 years*. With a lifetime of 10 years, this payback period would make the option viable and attractive. Hence, \$145/kWh is a good cutoff to achieve payback periods of 4 years. The payback periods are summarized in Table 5.2. In comparison, Battery A and B in the example were \$628/kWh and \$1526/kWh, respectively. Hence, there is potential to achieve reasonable payback periods with current battery technology provided the necessary incentives are present in the electricity markets using load shifting.

5.1.12 Additional Revenue Streams

Load shifting alone may not be enough to generate sufficient savings to warrant the capital expenditure. The number of cycles per year need to increase to generate additional revenue. One method is also to participate in market programs to generate revenues. Savings from participating in Frequency Regulation (FR) are 8–10 times the savings achieved by load shifting [44]. If the annual savings are increased by an order of magnitude, the payback period decreases by an order of magnitude, thus making these units even more financially attractive. When participating in market programs, batteries with high C-rates are more desirable since they have higher power output with the same capacity, thus they are able to provide greater flexibility for frequency regulation and generate more revenues.

5.1.13 Summary

A scalable control architecture for optimizing economics of large embedded battery applications is presented. The hierarchical structure permits the system-wide economic optimization to be solved in real-time through reduced computational complexity. The benefits of using electrical storage in conjunction with passive thermal energy storage are demonstrated via a simulation study of a 96-zone system. Load shifting is able to outperform the conventional tradeoff curve for constant setpoint control between comfort and cost minimization.

With the capital cost of manufacturing batteries decreasing, electrical storage is becoming a financially viable option alongside thermal energy storage. Using high performance large batteries does not give desirable payback periods for applications with load shifting only. However, lower performance smaller batteries have this potential. As battery costs decrease further over the foreseeable future, the return on investment becomes even more favorable. Supplementing load shifting savings

Table 5.3: Comparison of Central HVAC Plant and VRF Systems

	Central HVAC Plant	**VRF**
Cooling Resource	Chilled Water	Liquid Refrigerant
Heating Resource	Hot Water	Superheated Gas Refrigerant
Return Resource	Return Water	Gas/Liquid Refrigerant
Temperature Regulation Unit	AHU	Indoor VRF Unit
Heat Rejection Unit	Chillers/Boilers/RTUs	Outdoor VRF Unit
Airside Decision	Zone Temperature Setpoints or Discrete State of Equipment (on/off)	Zone Temperature Setpoints or Operating Mode (cooling/heating)
Waterside Decision	Equipment Operation Schedule	VRF Unit Operation Schedule

with revenues from programs such as frequency regulation help pay off the capital investment faster.

5.2 Variable Refrigerant Flow Systems

Variable refrigerant flow (VRF) is a type of HVAC technology that uses refrigerant as the cooling and heating fluid as opposed to the chilled water and hot water used in chiller plant and air handler unit (AHU)-based applications [5]. The VRF system consists of many indoor units (also known as fan-coil units) each serving an air space in buildings and several outdoor units (condensing units) that reject heat back to the ambient. Refrigerant is used as the heat transfer fluid that is exchanged between indoor units and outdoor units, whereas central HVAC plants discussed earlier use chilled and hot water to transfer heat between indoor air-handling units for zones and outdoor chillers and boilers. Most of the arguments presented in this work naturally extend to handle VRF systems with a change of terminology as shown in Table 5.3.

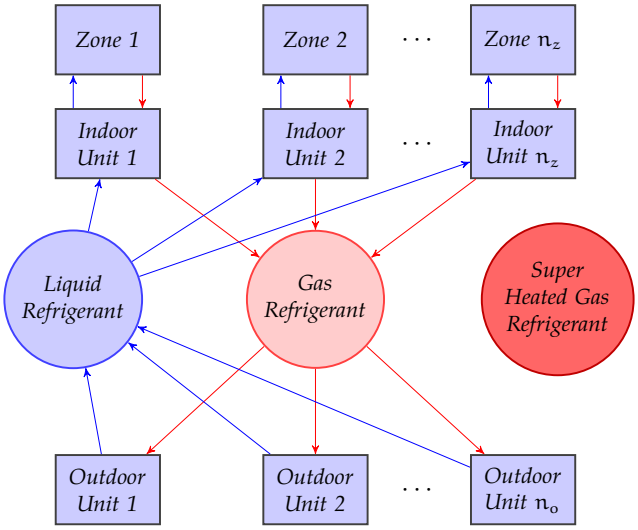

Figure 5.7: VRF system in cooling mode

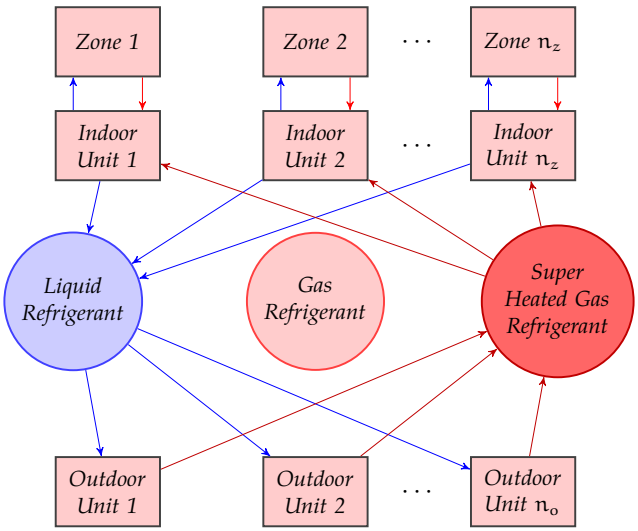

Figure 5.8: VRF system in heating mode

5.2.1 Modes of Operation

Various configurations exist for VRF systems. A schematic depicting the flow of refrigerants for one type of configuration is shown in Figure 5.7. A zone is defined to be the building space serviced by a single indoor VRF unit.

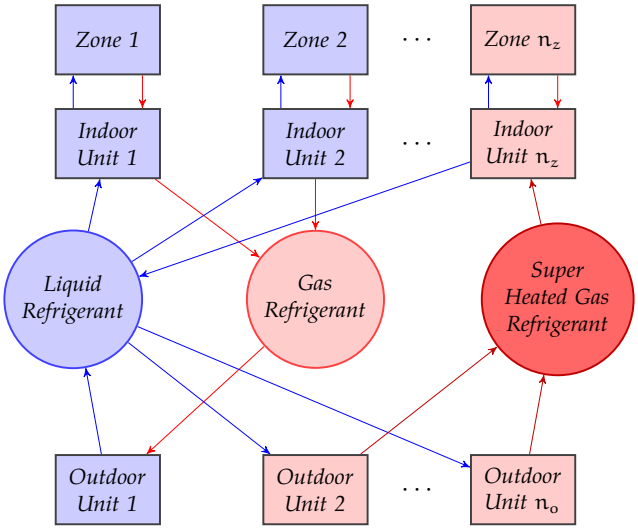

Figure 5.9: VRF system in mixed heating and cooling mode

A VRF system in heating mode is shown in Figure 5.8. A primary advantage of VRF systems is that some zone units may be in cooling mode while others are in heating mode simultaneously, as shown in Figure 5.9. This arrangements permits heat transfer between zones to improve comfort, which can be useful if some zones have glass exteriors where radiation places a large heat load and other zones are colder. A control system for VRF systems may have to make the cooling/heating mode decision for each zone as well.

5.2.2 VRF Systems Problem Statement

The goal for VRF systems is to develop a solution to solve the problem: given a forecast of electricity prices, ambient conditions, and disturbance loads, decide how to operate the VRF system to minimize the total operating cost:

- What setpoints are sent to the indoor unit zone controllers?
- When are outdoor units turned on, and at which load are they operating?

- When is storage being charged and when is it being discharged?

In making these decisions, the following constraints must be respected:

- The zone temperatures must lie within a predefined comfort region.
- Capacity constraints on equipment and storage tanks cannot be violated.
- Equipment cannot be switched on and off too rapidly.

With this problem statement, the same key ideas from Chapter 3 can be applied to address this class of systems, including VRF systems with a large number of zones. The primary change occurs in the cooling duty constraint, which is explored further in the next section.

5.2.3 Cooling Duty Constraint

In the general framework, the cooling duty model ($\dot{Q}_{c,i} = f(T_i, u_i)$) depends on the zone setpoints, u_i, that can be sent to the regulatory controllers associated with each indoor VRF units. Below, we discuss two specific instances for zone setpoints that are common in VRF systems: (1) cooling mode/flows and (2) temperature setpoints.

Cooling Mode/Flows. One potential case is that the supervisory MPC layer determines the mode (heating or cooling) for indoor VRF units along with the necessary refrigerant flow (cooling load). Let the binary variable z_i denote the mode of unit i where

$$z_i = \begin{cases} 1 & \text{Cooling} \\ 0 & \text{Heating (Not Cooling)} \end{cases}$$

Let m_i denote the cooling refrigerant flow to unit i. With $u_i = \begin{bmatrix} z_i \\ m_i \end{bmatrix}$, a simplified cooling duty model can be derived

$$\dot{Q}_{c,i} = \begin{cases} \alpha_i(T_i) \cdot m_i & z_i = 1 \\ 0 & z_i = 0 \end{cases}$$

in which $\alpha_i(T_i)$ is a conversion factor between cooling duty and refrigerant flow that depends on the zone temperature. If $\alpha_i(T_i)$ is a constant, the low-level problem with this cooling duty model can be formulated as a mixed-integer linear program (MILP).

Temperature Setpoint. Another case is the supervisory MPC layer sends temperature setpoints to each zone indoor unit. With $u_i = T_{sp,i}$, the formulation presented in Chapter 3 is also applicable to VRF systems. The main difference occurs in the numerical model parameters used to predict the cooling duty $\dot{Q}_{c,i}$ from the temperature setpoint and zone temperature.

5.2.4 Summary

In both cases, the ideas presented for the central HVAC plant case naturally extend to address VRF systems as well. Different solvers may be required if the low-level problem becomes a mixed-integer optimization, but the decomposition framework is made general to increase the range of its applicability.

6 CASE STUDY

In this chapter, we present a large-scale industrially relevant case study where solving a single MPC optimization problem is not feasible for real-time implementations. The study is loosely based on the Stanford University campus, consisting of both an airside and waterside system. The airside system includes 500 zones spread throughout 25 campus buildings, and the waterside system includes the central plant equipment, such as chillers, used to meet the load from the airside zones. Active TES is also included. The models and parameters from this case study are made publicly available for other researchers interested in designing control strategies for managing these applications. This case study has also been presented in [72, 74].

6.1 Motivation

The hierarchical decomposition proposed in Chapter 3 is one way to decompose the combined problem. However, there may be alternative methods as well. Several MPC-based control schemes have been proposed in the literature (see [57, 59, 96] as examples). It may become difficult to evaluate these novel decomposition strategies without having a common system against which to benchmark performance. Additionally, large-scale test cases are often not available so it may not be possible to demonstrate scalability of these approaches. To address these needs, a case study problem definition based on an industrial application has been constructed. This case study is made publicly available for other researchers in the HVAC community to design and test the performance and viability of various control systems. The aim of the release is to provide a standardized industrially relevant problem for the research community.

Figure 6.1: The new Stanford campus central plant constructed as part of the $485-million Stanford Energy System Innovations (SESI) project [9]

6.2 Stanford University Central Plant

The case study is inspired by the Stanford University campus [84]. Stanford University has recently replaced an aging natural gas cogeneration plant with a new heat-recovery system to service the cooling and heating loads of their campus as part of the $485-million Stanford Energy System Innovations (SESI) project [9]. In addition to the HRCs to improve overall efficiency, hot and cold thermal energy storage tanks were constructed. These large insulated tanks, along with the rest of the central HVAC plant, are depicted in Figure 6.1. Johnson Controls designed and commissioned the control architecture for the new central plant. Results have shown that the MPC-based system achieves 10–15% more energy cost savings compared to the best team of trained human operators [92]. While this project was focused primarily on optimization of the waterside, the case study is being extended to include treatment of the airside system as well. A simplified version of the Stanford project is used to highlight the complexity of controlling a large-scale combined

airside and waterside system while removing some of the problem features and intricate details to increase clarity for a research perspective.

6.3 System Definition

The HVAC system for the case study contains of a central plant that provides the cooling needs of a 500-zone campus. The central plant has eight conventional chillers as well as the supporting pumps and cooling towers for these chillers. Heating equipment, such as boilers or heat-recovery chillers, is omitted to simplify the problem. Each chiller has minimum and maximum cooling capacities of 2.5 MW and 12.5 MW, which gives a total central plant capacity of 100 MW cooling. In addition to the passive thermal energy storage present in the form of building mass, there is active thermal energy storage with a chilled water tank. This active TES tank has a maximum capacity of 100 MWh cooling.

The 500-zone campus consists of 25 buildings, each with 20 zones that have independent local temperature controllers. All zone temperatures need to be kept between 20.5 and 22.5 °C to ensure occupant comfort 24 hours a day, 7 days a week. The airside models describe the temperature dynamics in each of the 500 zones, and the waterside models describe the power consumption of the central plant equipment. These models for the equipment and zones are given in Section 6.5.

The control system must determine the zone temperature setpoints and waterside equipment operation schedule. The economic objective is to minimize energy costs in the presence of time-varying electricity prices and a peak demand charge as well as environmental disturbances such as weather while meeting constraints on comfort and equipment.

Table 6.1: Parameters

Variable	Description	Data Field	Unit
t	Time	`param.time`	h
T_a	Ambient Temperature	`param.AmbientTemp`	°C
c_k	Electricity Pricing	`param.ElecPrices`	US$
c_{peak}	Monthly Demand Charge	`param.DemandCharge`	US$
T_{min}	Lower Bound of Comfort Zone	`param.ComfortMin`	°C
T_{max}	Upper Bound of Comfort Zone	`param.ComfortMax`	°C
s_{max}	Active TES Capacity	`param.StorageCapacity`	kWh

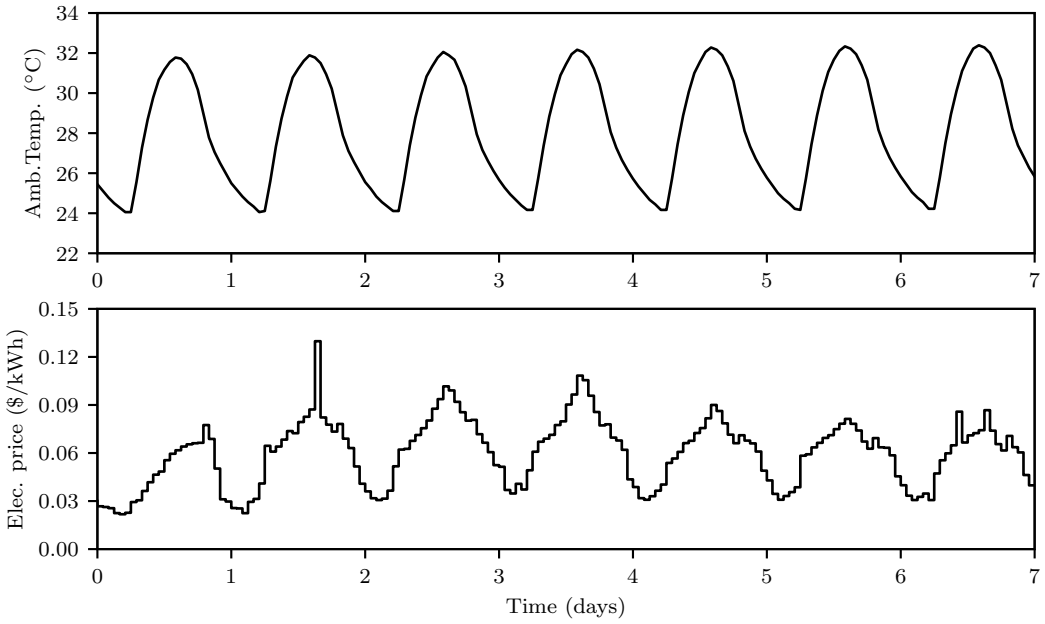

Figure 6.2: Weather and pricing data for the case study

6.4 Parameters

The parameters for this case study are summarized in Table 6.1. The data field refers to the location of this parameter in the data file released on the case study website.

The main disturbance considered in this case study is the ambient temperature. Typical ambient temperature data during the summer for a city in the southern United States is shown in Figure 6.2. To reject the loads back to the ambient, the

HVAC system buys power from the electricity market and runs the central plant equipment. Two components of the pricing structure are considered in this study: time-of-use charges, which assess time-varying prices on electricity use throughout the day, and peak demand charges, which are proportional to maximum rate of power consumption over period of time (typically a month). Electricity pricing data obtained from Johnson Controls over a week-long period is given in Figure 6.2. The monthly peak demand charge for this market is \$4.56/kW.

6.5 Models

Next, the models used for the airside and waterside system are presented.

6.5.1 Airside System

In the airside system, models are used for zone temperature dynamics. The dynamics of cooling a single zone or building can be represented by an energy balance as stated in Chapter 2. The lumped model for the temperature of zone i as given by

$$C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{other,i} \tag{6.1}$$

in which C_i is the thermal capacitance of the zone, H_i is a scaled heat transfer coefficient with the ambient, T_a is the ambient temperature, $\dot{Q}_{c,i}$ is the cooling rate from the HVAC system, $\dot{Q}_{other,i}$ is an external load place on the zone, and β_{ij} characterizes the degree of coupling between zones i and j. For non-adjacent zones, $\beta_{ij} = 0$. The heat transfer is depicted in Figure 6.3.

The supervisory control system determines the zone temperature setpoints. Using an ideal PI controller, the linear cooling duty controller model relating the zone temperature setpoint $T_{sp,i}$ to the cooling rate $\dot{Q}_{c,i}$ delivered to the zone is given

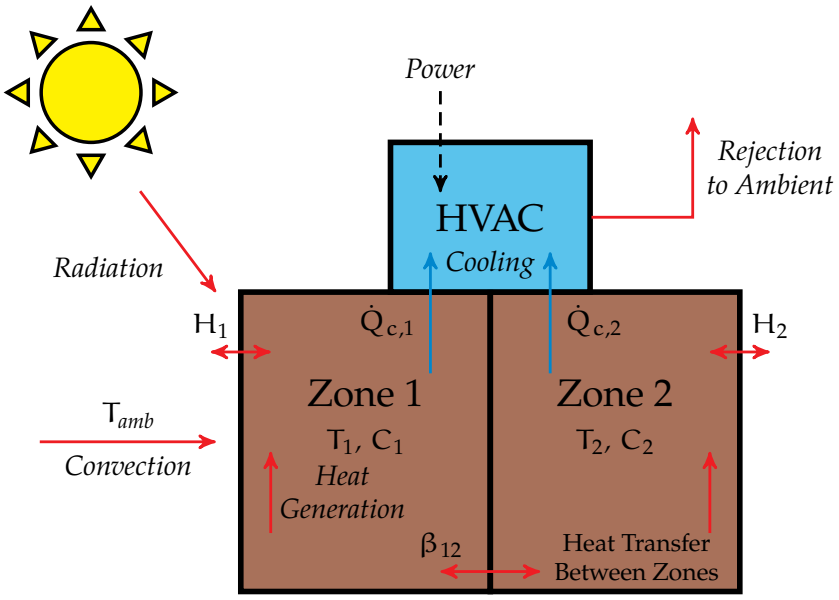

Figure 6.3: Diagram of airside heat transfer.

by

$$\dot{Q}_{c,i} = \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right]$$
$$\varepsilon_i = T_{sp,i} - T_i \tag{6.2}$$

in which $K_{c,i}$ and $\tau_{I,i}$ are the PI controller parameters and ε_i is the tracking error. The airside model parameters for this case study are provided in the associated data file and are summarized in Table 6.2. In the data file that is released, these airside parameters are grouped by building and presented as a cell array of length of 25 with each cell containing the parameter values for the 20 zones in that particular building.

Table 6.2: Airside Model Parameters

Variable	Description	Data Field	Unit
H_i	Scaled Heat Transfer Coefficient with Ambient	`airside.H`	kW/K
C_i	Thermal Capacitance	`airside.C`	kJ/K
β_{ij}	Scaled Heat Transfer Coefficient between Zones	`airside.Beta`	kW/K
$\dot{Q}_{ss,i}$	PI Steady State Cooling	`airside.Qss`	kW
$K_{c,i}$	PI Controller Gain	`airside.Kc`	kW/K
$\tau_{I,i}$	PI Integral Time Constant	`airside.tauI`	h

6.5.2 Waterside System

In the waterside system, models are used for central plant equipment electricity consumption and storage tank dynamics. The equipment models are static, and the storage model is dynamic. Static equipment models provide resource consumption as a function of relevant inputs for a given steady-state operating point. The effects of transient dynamics during equipment startup and shutdown are handled by local regulatory controllers. Frequent startups and shutdowns are prevented by enforcing explicit dwell time constraints.

However, storage tank models are necessarily dynamic, since they are used for load shifting. It is assumed throughout this section that the enthalpy H of a volume V of water at uniform temperature T is given by

$$H = \rho C_p V (T - T_{ref}) \tag{6.3}$$

in which ρ is the (constant) density, C_p is the (constant) heat capacity, and T_0 is an arbitrary reference temperature.

For the central chilling plant in the case study, the three types of equipment are chillers, cooling towers, and pumps. Each chiller is modeled using the semi-empirical

Gordon-Ng model, [50] according to

$$W_{CH} \coloneqq \left(Q_{CH} + a_1 T_{CHWS} + a_2 \left(1 - \frac{T_{CHWS}}{T_{CWS}} \right) \right) \frac{T_{CWS}}{T_{CHWS} - a_3 Q_{CH}} - Q_{CH} \tag{6.4}$$

in which the parameters a_1, a_2, and a_3 are obtained via regression with measured equipment data. The temperatures are assumed to be fixed parameters in this study. The volumetric flow can then be calculated from the enthalpy relationship in Equation (6.3). Each cooling tower is modeled with a simplified effectiveness model [36] for calculating cooling duty, with a simple cubic fit for fan electricity [10].

$$Q_{CT} = Q_{CH} + Q_{CH} \coloneqq \frac{c_1 (m_{CW})^{c_3}}{1 + c_2 \left(\frac{m_{CW}}{m_{air}} \right)^{c_3}} (T_{CWR} - T_{WB}) \tag{6.5}$$

$$W_{CT} \coloneqq \kappa (m_{air})^3 \tag{6.6}$$

At fixed T_{CWR} and T_{CWS}, the water flow rate m_{CW} is calculated from an enthalpy balance. Then, using the known T_{WB}, (6.5) can be rearranged to solve for the required m_{air}, which is then used in (6.6) for electricity calculation. Coefficients c_1, c_2, c_3, and κ are obtained via regression. Finally, pumps are modeled with a black-box empirical model

$$W_P \coloneqq b_1 \ln (1 + b_2 V_{CHW}) + b_3 V_{CHW} + b_4 \tag{6.7}$$

with regression coefficients b_1 through b_4. Note that the flows V_{CW} and m_{CW} are obtained from Q_{CH} and Q_{CT} via the appropriate constant-heat-capacity energy balances.

Active storage tanks are modeled using a two-layer stratified tank model similar to Ma et al. [54]. The hot and cold sections in the tank are each assumed to be uniform in temperature, with heat exchange between the two layers proportional

to their temperature difference. The tank is well-insulated, but heat exchange takes place between the two layers. With state vector defined as $(V_{cold}, V_{hot}, H_{cold}, H_{hot})$, the differential equations are expressed as follows:

$$\begin{aligned}
\frac{dV_{hot}}{dt} &= -\nu_+ + \nu_-, \\
\frac{dV_{cold}}{dt} &= \nu_+ - \nu_-, \\
\frac{dH_{hot}}{dt} &= -\frac{H_{hot}}{V_{hot}}\nu_+ + h_-\nu_- - K\left(\frac{H_{hot}}{V_{hot}} - \frac{H_{cold}}{V_{cold}}\right), \\
\frac{dH_{cold}}{dt} &= h_+\nu_+ - \frac{H_{cold}}{V_{cold}}\nu_- + K\left(\frac{H_{hot}}{V_{hot}} - \frac{H_{cold}}{V_{cold}}\right),
\end{aligned}$$

in which inputs ν_+ and ν_- are the charge and discharge volumetric flows, and parameters h_+ and h_- are the (per-volume) supply and return enthalpies, with heat transfer coefficient K between the two layers. The corresponding temperatures T_{hot} and T_{cold} are calculated from the linear enthalpy relationship (6.3), as are the volumetric enthalpies h_+ and h_- of the incoming streams. In chilled water tanks, the state of interest is the enthalpy of the cold section H_{cold}, while for hot water tanks, it is H_{hot}. Total volume $V_{hot} + V_{cold} = V_{total}$ is held constant.

For optimization, a linear approximation of the nonlinear tank model is used

$$\frac{ds}{dt} = -\sigma s + \eta \dot{Q}_{storage} \tag{6.8}$$

in which $s \coloneqq H_{cold}$ is the enthalpy of the cold section and $\dot{Q}_{storage}$ is the rate of cold enthalpy inflow (positive) or outflow (negative). The coefficients σ and η are identified from data.

All waterside parameters are shown in Table 6.3, with numerical values given in the data file.

Table 6.3: Waterside Model Parameters

Variable	Description	Data Field	Unit
T_{CHWS}	Constant chilled water supply temperature	`waterside.common.Tchws`	°C
T_{CHWR}	Constant chilled water return temperature	`waterside.common.Tchwr`	°C
T_{CWS}	Constant chilled water supply temperature	`waterside.common.Tchwr`	°C
ρ	Water density	`waterside.common.rho`	kg/m^3
C_p	Water heat capacity	`waterside.common.Cp`	kJ/(kg °C)
T_{ref}	Enthalpy reference temperature	`waterside.common.Tref`	°C
Q_{CH}^{min}	Minimum chiller cooling duty	`waterside.chiller.Qmin`	kW
Q_{CH}^{max}	Maximum chiller cooling duty	`waterside.chiller.Qmax`	kW
a_i	Chiller regression coefficients	`waterside.chiller.a`	varies
Q_{CT}^{min}	Minimum cooling tower duty	`waterside.tower.Qmin`	kW
Q_{CT}^{max}	Maximum cooling tower duty	`waterside.tower.Qmax`	kW
c_i	Cooling tower regression coefficients	`waterside.tower.c`	varies
κ	Cooling tower fan coefficient	`waterside.tower.kappa`	kJ/kg
V^{min}	Minimum pump flow rate	`waterside.pump.Vmin`	m^3/s
V^{max}	Maximum pump flow rate	`waterside.pump.Vmax`	m^3/s
b_i	Pump regression parameters	`waterside.pump.b`	varies
V_{total}	Total tank volume	`waterside.tank.Vtot`	m^3
K	Layer heat transfer coefficient	`waterside.tank.K`	m^3/s

6.6 Availability

The entire data set and model parameters for the case study are made publicly available for researchers in the HVAC community on the following website: `https://hvacstudy.github.io/`. An image of the main page for this website is shown in Figure 6.4. In addition to the data file containing the parameters and models presented earlier, documentation is also provided outlining model forms and additional background information.

6.7 Sample Results

6.7.1 Control Architecture

With the case study system defined, a supervisory control system can be used to make decisions. We consider the two-layer hierarchical control system for large-

Figure 6.4: Website for case study release: `https://hvacstudy.github.io/`

scale application presented in Chapter 3. In this control architecture, the nonlinear equipment models in (6.4) through (6.7) are approximated as piecewise-linear. This choice permits the low-level waterside problem to be formulated as an MILP, and arbitrary accuracy can be achieved by this approximation by using more pieces, although at the cost of computational speed [86]. The linear model for the storage tank is used instead of the full nonlinear model.

6.7.2 Optimization Results

This decomposition was applied to the case study system. The results from the high-level optimization are shown in Figure 6.5. The top plot shows the aggregate load for the campus and how it is met using storage and production from equipment. The bottom plot shows the 25 average building temperature profiles. A combination of both active and passive thermal energy storage is used to shift the power load from

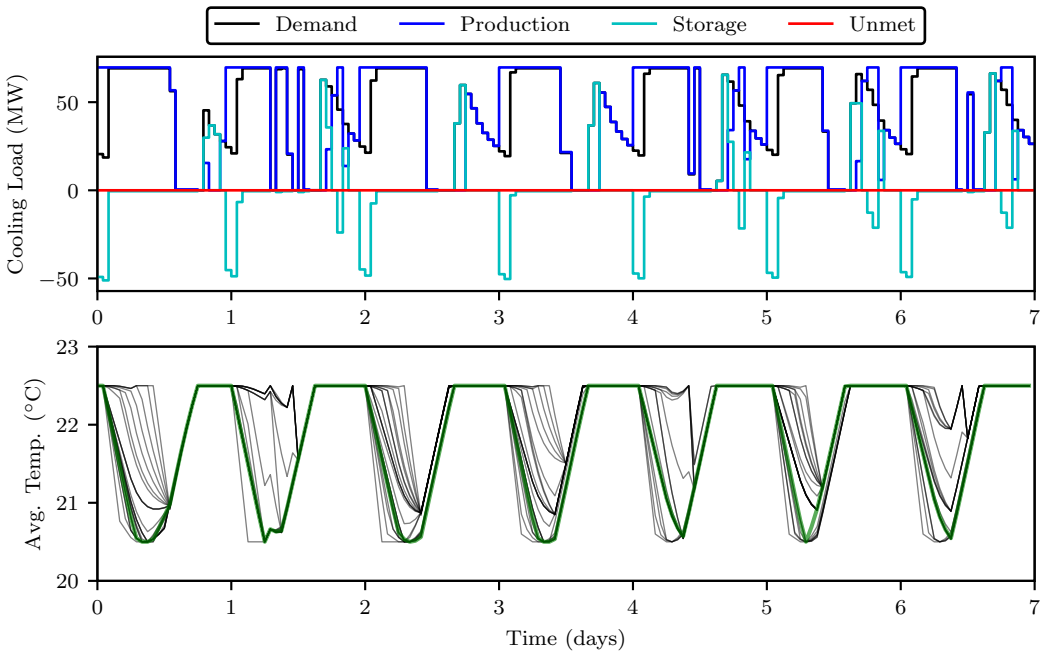

Figure 6.5: Results from the high-level optimization for the case study

peak hours to off-peak hours. Since the active TES tank is not very large, passive TES is utilized by precooling the buildings. These load profiles are sent to the low-level problems.

Figure 6.6 shows the low-level airside results for all 500 zones. The top plot shows the zone temperature profiles, and the bottom plot shows the zone temperature setpoints required to obtain those temperatures. In each building, the order of precooling is chosen to minimize energy usage and is determined by the physical characteristics of the zones. As noted previously, the initial spikes in setpoints during the precooling stages are used to overcome sluggish regulatory responses.

Figure 6.7 shows the equipment operation schedule computed by the low-level waterside optimization to meet the load from the high-level problem. The top plot shows the updated total production plot with the more detailed equipment models. The bottom plot is the Gantt chart depicting when equipment is turned on and off

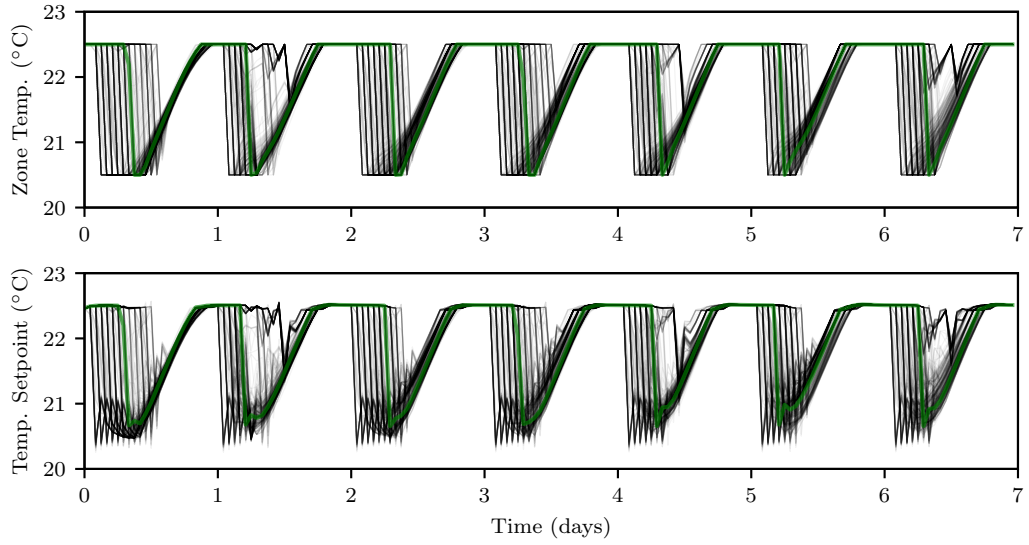

Figure 6.6: Results from the low-level airside optimization for the case study

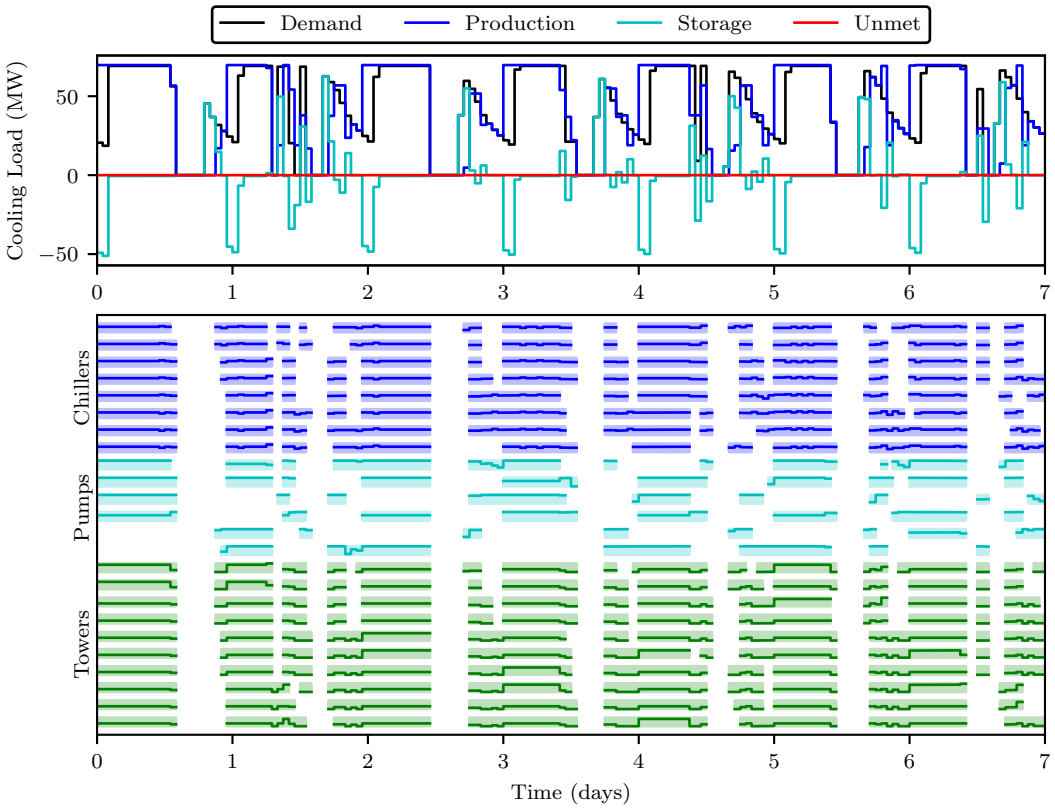

Figure 6.7: Results from the low-level waterside optimization for the case study

and at which load it is running. The total cost for this control policy is $78,700. For comparison, the cost without load shifting is $94,900, hence using optimization with TES results in 17% cost savings.

6.7.3 Real-Time Computational Requirements

The optimization problems were solved using Gurobi 7.5.2 via MATLAB R2017b on a desktop machine with 16GB RAM and 2.70GHz Intel® Core™ Processor i5-7500T. The high-level problem was solved in 1.3 seconds. Each of the 25 low-level problems was solved in approximately 1.2 seconds. The low-level waterside subproblem was given two minutes of computation time, after which the best feasible solution (with an optimality gap of 0.2%) was accepted. Since the low-level problems can be solved in parallel, the total computational time required in practice is about 2 min. Since control executions occur every 15 minutes, this framework can be deployed online for real-time applications.

6.8 Summary

A case study partly based on the Stanford University campus HVAC plant has been presented. The models and data from this study are made publicly available for HVAC researchers online at `https://hvacstudy.github.io/`. The proposed decomposition has been demonstrated on this benchmark system, achieving 17% energy cost savings over the baseline. Additionally, the decomposition can be implemented online due to its low computational requirements.

7 ZONE TRACKING MPC

The primary focus of this dissertation is the design of economic MPC-based control architectures for building systems. In these formulations, comfort bounds on temperatures are enforced as soft constraints in the MPC formulations. However, while desirable, these bounds do not needed to be satisfied exactly at all times as physical constraints on capacity need to be. When a large disturbance hits the building, the HVAC system may not be able to reject the heat to the ambient and bring the temperature back inside the comfort region immediately. It may take several controller time steps to do so. In this chapter, an alternate method of controlling such systems, called *zone tracking MPC*, is proposed [66].

First, the conventional setpoint tracking MPC paradigm is discussed. Then, two motivating examples for which setpoint tracking is not sufficient are given. The zone tracking MPC problem is outlined, and a nominal stability proof is provided. Simulations are performed on the two motivating examples to demonstrate the effectiveness of the zone tracking MPC scheme.

7.1 Conventional Setpoint Tracking MPC

The conventional MPC paradigm considers tracking to a setpoint [83]. A dynamic system model is used to make predictions of the output variables as a function of the actions taken by the controller. An online optimization problem is solved over a horizon to drive process variables to desired setpoints.

Let the system model be represented as

$$x(k+1) = f(x(k), u(k)) \tag{7.1}$$

For ease of notation, this model is referred to as $x^+ = f(x, u)$ with the time index dropped and x^+ used to denote the value of the state at the next time step. The state and input are subject to constraints

$$x(k) \in \mathbb{X}, \qquad u(k) \in \mathbb{U} \qquad \text{for all } k \in \mathbb{I}_{\geqslant 0}$$

Given a finite horizon (integer N) and an input sequence $\mathbf{u}$ of length N, $\mathbf{u} = \{u(0), u(1), \ldots, u(N-1)\}$, let $\phi(k; x, \mathbf{u})$ denote the solution of (7.1) at time k for a given initial state $x(0) = x$. The terminal region constraint is

$$\phi(N; x, \mathbf{u}) \in \mathbb{X}_f \subseteq \mathbb{X}$$

The set of feasible initial states and associated control sequences are

$$\mathbb{Z}_N = \{(x, \mathbf{u}) \mid u(k) \in \mathbb{U}, \phi(k; x, \mathbf{u}) \in \mathbb{X} \text{ for all } k \in \mathbb{I}_{0:N-1}, \text{and } \phi(N; x, \mathbf{u}) \in \mathbb{X}_f\}$$

The set of feasible initial states is

$$\mathcal{X}_N = \{x \in \mathbb{R}^n \mid \exists \mathbf{u} \in \mathbb{U}^N \text{ such that } (x, \mathbf{u}) \in \mathbb{Z}_N\}$$

For each $x \in \mathcal{X}_N$, the corresponding set of feasible input sequences is

$$\mathcal{U}_N = \{\mathbf{u} \mid (x, \mathbf{u}) \in \mathbb{Z}_N\}$$

For any state $x \in \mathbb{R}^n$ and input sequence $\mathbf{u} \in \mathbb{U}^N$, define the objective as

$$V_N(x, \mathbf{u}) = \sum_{k=0}^{N-1} l(\phi(k; x, \mathbf{u}), u(k)) + V_f(\phi(N; x, \mathbf{u}))$$

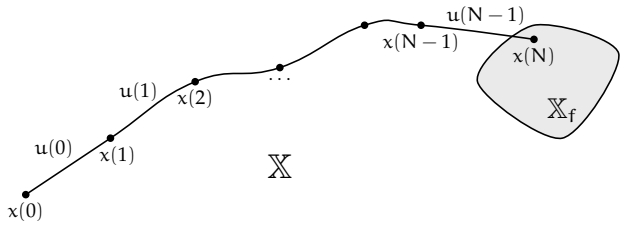

Figure 7.1: Setpoint tracking MPC with a terminal region

in which $\ell(x, u)$ is the stage cost and $V_f(x(N))$ is the terminal cost. Quadratic costs are typically used when tracking to the origin and minimizing input usage with

$$\ell(x, u) = x^T Q x + u^T R u$$
$$V_f(x(N)) = x(N)^T P_f x(N)$$

Deviation variables can be used when the setpoint is not the origin to obtain costs of the form $(x - x_{sp})^T Q (x - x_{sp})$ in which x_{sp} is the setpoint.

With these definitions, the finite horizon optimal control problem (OCP) is formulated as

$$\begin{aligned} \text{OCP} : \min_{\mathbf{u} \in \mathcal{U}_N} V_N(x, \mathbf{u}) = \min \quad & \sum_{k=0}^{N-1} l(x(k), u(k)) + V_f(x(N)) \\ \text{s.t.} \quad & x^+ = f(x, u) \\ & u \in \mathbb{U}, x \in \mathbb{X} \\ & x(N) \in \mathbb{X}_f \subseteq \mathbb{X} \end{aligned} \tag{7.2}$$

One feasible trajectory for this problem is depicted in Figure 7.1. The optimization problem is solved to obtain the solution sequence $\mathbf{u}^0$. The first control input is injected into the system. This control law is denoted as

$$\kappa_N(x) = u^0(0; x) \tag{7.3}$$

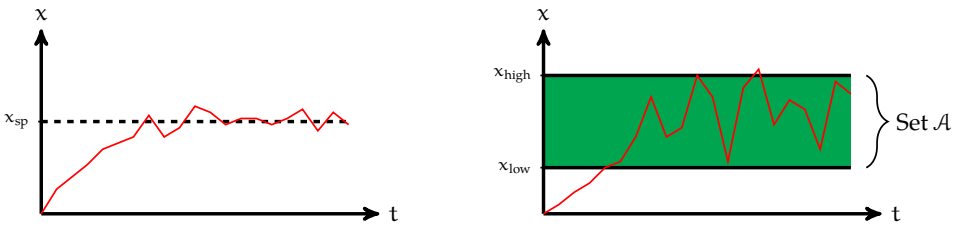

(a) Setpoint Tracking (b) Zone Tracking

Figure 7.2: Comparison of setpoint vs zone tracking formulations

The optimum may not necessarily be unique, in which case $\kappa_N(x)$ is a point-to-set map. At the next controller time step, the optimization problem is resolved using an updated measurement of the state as $x(0)$. This paradigm has been successful in practice with thousands of applications in the chemical and petroleum industries alone in the past two decades [78].

7.2 Motivation for Zone Tracking

While the setpoint tracking formulation addresses many systems, the control objective in several key industrial applications is to keep process variables within a region, or zone, as shown in Figure 7.2. For example, in a buffer tank between upstream and downstream process units, the goal is to reject upstream flow disturbances so that they do not propagate downstream. The height of material in the tank is not tracked to a setpoint, but it should be kept within physical bounds to prevent emptying or overflowing.

Another example occurs in commercial building temperature control. The aim is to maintain the air temperature between bounds to keep the occupants in the building comfortable. If the temperature is outside of this zone, the controller should take action to track this "comfort" zone.

One approach is to make the tracking penalty, Q, small in the cost functions to

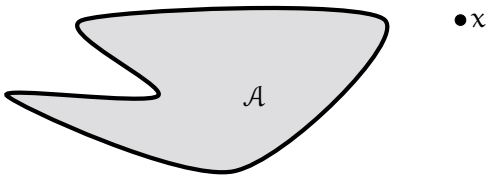

Figure 7.3: Target set for zone tracking

permit excursions away from the setpoint in order to minimize input usage. This method is called *loose setpoint tracking*. However, as shown later in this chapter, it can lead to undesirable consequences when large disturbances enter the system.

For such applications, the aim is to apply model predictive control (MPC) for optimal zone tracking. Consider tracking the state to the target set $\mathcal{A}$ in Figure 7.3. When the state is outside of the zone, the controller should take action to drive the system towards the zone. When the state is inside of the zone, the controller should maintain the system there. To achieve the objective of tracking a closed set $\mathcal{A}$, the cost function should be 0 if $x \in \mathcal{A}$, otherwise it should be positive. The distance function is one such function that satisfies this property

$$\ell(x, u) = |x|_{\mathcal{A}} = \inf_{z \in \mathcal{A}} |x - z| \tag{7.4}$$

Since the traditional MPC paradigm involves tracking to a setpoint, the theory needs to be extended to handle this case to ensure that asymptotic stability is achieved for an MPC controller with this type of cost function.

7.3 Background

Before posing the zone tracking problem, some remarks are made on existing approaches for these systems.

Several industrial software packages are available for the control of these appli-

cations. Aspen Technology provides upper and lower soft constraint capabilities with their software Aspen DMCplus [78]. Honeywell uses funnel constraints in their software. However, the theoretical conditions under which the closed-loop system is stable are not well understood.

MPC for tracking zone regions has been previously proposed and applied to a distillation column [25]. However when process variables are inside the target zone though, it may be desirable to achieve secondary economic objectives. In the buffer tank example, since the aim is to minimize the propagation of the disturbance downstream, the controller should minimize the rate of change of the downstream flowrate whenever possible. In the building temperature control problem, if the temperature is inside the comfort zone, the controller should ideally minimize the energy usage or costs. Incorporating such secondary economic objectives can provide additionally flexibility over tracking to a desirable operating point such as the middle of the zone.

When economic objectives are desired, economic MPC (EMPC) is typically utilized [16, 79]. A multiobjective approach to economic MPC with a stabilizing tracking constraint has been proposed but not extended to zone tracking [100]. One variant of EMPC is Lyapunov-based economic MPC, which keep the state within a zone region defined by a sublevel set of the Lyapunov function [22]. However, since the region is a sublevel set, the zone cannot be an arbitrary design choice by the user.

Furthermore, actuators may be discrete (e.g., valves may be on-off). The zone tracking problem formulation should consider the discrete actuator case as well.

In this work, the framework is extended from setpoint tracking to zone tracking and the conditions under which the closed-loop system is stable are stated. The zone is made to be a design choice and secondary economic objectives can be achieved when inside the zone. Discrete actuators are also addressed.

7.4 Problem Setup

The setup of the zone tracking problem $\mathbb{P}_N(x)$ is similar to the conventional MPC problem in Equation (7.2) with the set $\mathcal{A}$ replacing the terminal region $\mathbb{X}_f$.

$$
\begin{aligned}
\mathbb{P}_N(x) : \min_{\mathbf{u}} \quad & \sum_{k=0}^{N-1} \ell(x(k), u(k)) && \text{Objective} \\
\text{s.t.} \quad & x^+ = f(x, u) && \text{Dynamic System Model} \\
& u \in \mathbb{U}, x \in \mathbb{X} && \text{Input and State Constraints} \\
& x(N) \in \mathcal{A} \subseteq \mathbb{X} && \text{Terminal Region}
\end{aligned}
$$

Without a terminal region, a terminal cost function $V_f(x(N))$ is no longer required. This zone tracking generalization encompasses the setpoint tracking problem with a terminal constraint as a special case when $\mathcal{A}$ is a point. The stage cost $\ell(x, u)$ is chosen as prescribed earlier: $\ell(x, u) = 0$ when $x \in \mathcal{A}$ and $\ell(x, u) > 0$ when $x \notin \mathcal{A}$.

7.5 Stability Theory

Assumption 7.1. *[Control Invariance]. $\mathcal{A}$ is a closed set with $\mathcal{A} \subseteq \mathbb{X} \subset \mathbb{R}^n$. For each $x \in \mathcal{A}$, there exists a $u \in \mathbb{U}$, denoted $\kappa_f(x)$, such that $f(x, u) \in \mathcal{A}$. The stage cost, $\ell(x, u)$, is 0 if $x \in \mathcal{A}$ and greater than 0 if $x \notin \mathcal{A}$.*

Assumption 7.2. *[System and Cost Continuity]. The functions $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ and $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}_{\geqslant 0}$ are continuous.*

Assumption 7.3. *[Optimal Cost Continuity]. The optimal cost $V_N^0(x)$ is continuous at the boundary of $\mathcal{A}$.*

Assumption 7.4. *[Constraint Set Properties]. The set $\mathbb{U}$ is compact. The set $\mathbb{X}$ is closed.*

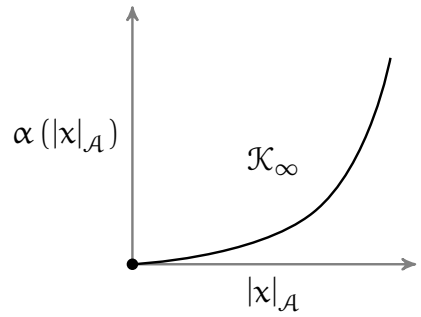

Figure 7.4: Example of a $\mathcal{K}_\infty$ function, $\alpha(\cdot)$.

Assumption 7.5. *[Cost Bound]. The stage cost satisfies*

$$\ell(x, u) \geqslant \alpha(|x|_{\mathcal{A}}) \quad \forall x \in \mathcal{X}_N, \forall u \in \mathbb{U}$$

in which $\alpha(\cdot)$ is a $\mathcal{K}_\infty$ function.

Recall that a $\mathcal{K}_\infty$ function $\alpha(\cdot) : \mathbb{R} \to \mathbb{R}_{\geqslant 0}$ is continuous, zero at zero, strictly increasing, and is unbounded, as depicted in Figure 7.4 [39].

Definition 7.6. *[MPC Control Law]. Let $u^0(0; x) := \kappa(x)$ denote the control law obtained by solving the MPC optimization problem.*

Theorem 7.7. *[Nominal Stability Result]. Under Assumptions 7.1 through 7.5, the set $\mathcal{A}$ is asymptotically stable for the closed-loop system $x^+ = f(x, \kappa(x))$ with a region of attraction of $\mathcal{X}_N$.*

Proof of Theorem 7.7

The proof involves showing that the the optimal cost $V_N^0(\cdot)$ is a Lyapunov function for the closed-loop system, which is a similar procedure for the nominal setpoint tracking MPC problem [83, Chapter 2]. To show that $V_N^0(\cdot)$ is a Lyapunov function

for the closed-loop system, the following three properties are required for all $x \in \mathcal{X}_N$

$$\begin{aligned} V_N^0(x) &\geqslant \alpha_1(|x|_{\mathcal{A}}) && \text{Lower Bound} \\ V_N^0(x) &\leqslant \alpha_2(|x|_{\mathcal{A}}) && \text{Upper Bound} \\ V_N^0(f(x, \kappa_N(x))) &\leqslant V_N^0(x) - \alpha_3(|x|_{\mathcal{A}}) && \text{Cost Decrease} \end{aligned}$$

in which $\alpha_1(\cdot), \alpha_2(\cdot), \alpha_3(\cdot)$ are $\mathcal{K}_\infty$ functions.

The control invariance of the target set is first shown followed by the proof of these three Lyapunov function properties.

Control Invariance. Assumption 7.1 ensures that the set $\mathcal{A}$ is control invariant. Suppose $x \in \mathcal{A}$, then there exists $u_1 \in \mathbb{U}$ such that $f(x, u_1) \in \mathcal{A}$. Since $f(x, u_1) \in \mathcal{A}$, there exists $u_2 \in \mathbb{U}$ such that $f(f(x, u_1), u_2) \in \mathcal{A}$. Continue the recursion for any N, and construct a sequence of the resulting inputs as

$$\tilde{\mathbf{u}}(x) = \{u_1(x), u_2(x), \ldots, u_{N-1}(x), u_N(x)\}$$

The corresponding state trajectory lies within the set $\mathcal{A}$

$$\phi(k; x, \tilde{\mathbf{u}}) \in \mathcal{A} \qquad \text{for all } k \in \mathbb{I}_{\geqslant 0}, k \leqslant N$$

Since the trajectory lies within $\mathcal{A}$, the total cost for this input sequence is

$$V_N(x, \tilde{\mathbf{u}}) = 0$$

based on the properties of the stage cost stated in Assumption 7.1. The optimal cost cannot be worse

$$V_N^0(x) \leqslant V_N(x, \tilde{\mathbf{u}}) = 0 \tag{7.5}$$

Suppose that the set $\mathcal{A}$ is not control invariant under MPC. At a later time, $\phi(k; x, \mathbf{u}) \notin \mathcal{A}$. Based on Assumption 7.1, the cost of the resulting trajectory must be control

$$V_N^0(x) > 0$$

which contradicts the above condition. Hence, $\mathcal{A}$ must be control invariant.

Lower Bound. Given the definition of $V_N(x, \mathbf{u})$ as a sum of stage cost and using Assumptions 7.1 and 7.5 gives

$$V_N(x, \mathbf{u}) = \sum_{k=0}^{N-1} \ell(\phi(k; x, \mathbf{u}), u(k)) \geqslant \ell(x, u(0; x)) \geqslant \alpha_1 \left(|x|_{\mathcal{A}}\right)$$

for all $x \in \mathcal{X}_N, \mathbf{u} \in \mathbb{U}^N$. Since it holds for all $\mathbf{u} \in \mathbb{U}^N$,

$$V_N^0(x) \geqslant \alpha_1 \left(|x|_{\mathcal{A}}\right) \quad \text{for all} \quad x \in \mathcal{X}_N$$

so the first property is established.

Upper Bound. Using Equation 7.5 and the properties of $\mathcal{K}_\infty$ functions yields

$$V_N^0(x) \leqslant 0 \leqslant \alpha_4 \left(|x|_{\mathcal{A}}\right) \quad \text{for all} \quad x \in \mathcal{A}$$

To establish the second property, this bound has to be extended from $\mathcal{A}$ to $\mathcal{X}_N$. From Assumption 7.2, since f and ℓ are continuous, we have that $V_N(x, \mathbf{u})$ is continuous, hence locally bounded. Additionally, Assumption 7.3 ensures continuity of $V_N^0(x)$ at the boundary of $\mathcal{A}$. With these properties, $V_N^0(x)$ is locally bounded on $\mathcal{X}_N$ [83, Proposition 2.15]. Since a function can be upper bounded by a $\mathcal{K}_\infty$ function if it is

locally bounded [83, Proposition 2.16], there exists $\alpha_2(\cdot) \in \mathcal{K}_\infty$ such that

$$V_N^0(x) \leqslant \alpha_2\left(|x|_{\mathcal{A}}\right) \quad \text{for all} \quad x \in \mathcal{X}_N$$

Cost Decrease. At a given state $x \in \mathcal{X}_N$, consider an optimal sequence $\mathbf{u}^0(x) = \{u(0;x), u(1;x), \ldots, u(N-1;x)\}$, and generate a candidate sequence for the successor state, $x^+ = f(x, \kappa_N(x))$ using

$$\tilde{\mathbf{u}}(x) = \{u(1;x), u(2;x), \ldots, u(N-1;x), \kappa_f(x(N))\}$$

with $x(N) := \phi(N; x, \mathbf{u})$. This sequence is feasible for x^+ because $\mathcal{A}$ is control invariant under control law $\kappa_f(\cdot)$ based on Assumption 7.1.
The cost for the candidate sequence is

$$V_N(x^+, \tilde{\mathbf{u}}) = V_N^0(x) - \ell(x, u(0;x)) + \ell(x(N), \kappa_f(x(N)))$$

Since the stage cost is nonnegative by construction,

$$V_N(x^+, \tilde{\mathbf{u}}) \leqslant V_N^0(x) - \ell(x, u(0;x))$$

The optimal cost cannot be worse, thus

$$V_N^0(x^+) \leqslant V_N(x^+, \tilde{\mathbf{u}}) \leqslant V_N^0(x) - \ell(x, u(0;x))$$

Using Assumption 7.5 for the stage cost gives

$$V_N^0(x^+) \leqslant V_N^0(x) - \alpha_1\left(|x|_{\mathcal{A}}\right) \quad \text{for all} \quad x \in \mathcal{X}_N$$

Therefore, the third property is established using $\alpha_1(\cdot) = \alpha_3(\cdot)$

$$V_N^0(f(x, \kappa_N(x))) \leqslant V_N^0(x) - \alpha_3(|x|_{\mathcal{A}}) \quad \text{for all} \quad x \in \mathcal{X}_N$$

Asymptotic Stability. Since the optimal cost $V_N^0(\cdot)$ is a Lyapunov function for the closed-loop system $x^+ = f(x)$ and control invariant, closed set $\mathcal{A}$, then $\mathcal{A}$ is asymptotically stable for $x^+ = f(x)$ [83, Theorem 2.13].

7.6 Secondary Objectives

Thus far, nothing has been said about the uniqueness of the solution to $\mathbb{P}_N(x)$. In fact, there is often a set of input sequences that can maintain the state within $\mathcal{A}$. With a nonunique control law, there exists the opportunity to achieve additional secondary control objectives by modifying the cost function in the problem. Other terms can be added to minimize economic costs (c) or input usage (R) or input movement (S) to select the most desirable operating point within $\mathcal{A}$

$$\bar{\ell}(x, u) = c'u + \frac{1}{2}u'Ru + \frac{1}{2}\Delta u'S\Delta u$$

Clearly, this type of stage cost would violate Assumption 7.1. In this section, a methodology for achieving secondary objectives is proposed while maintaining the nominal stability result.

First, the solution set of the zone tracking problem $\mathbb{P}_N(x)$ is characterized as

$$\mathbb{U}^0(x) = \{\mathbf{u} \in \mathcal{U}_N \mid V_N(x, \mathbf{u}) = V_N^0(x)\}$$

Based on the above theory, choosing any $\mathbf{u} \in \mathbb{U}^0(x)$ is stabilizing for closed-loop

system. To select the most desirable input, a second optimization problem is solved

$$\min_{\mathbf{u}\in\mathbb{U}^0(x)} \bar{V}_N(x,\mathbf{u}) = \min_{\mathbf{u}\in\mathbb{U}^0(x)} \sum_{k=0}^{N-1} c'u + \frac{1}{2}u'Ru + \frac{1}{2}\Delta u'S\Delta u$$

which gives a unique solution if either $R > 0$, $S > 0$, or the economic condition yields a nondenegerate optimal operating point. Based on the control objectives of the application, some of these terms can be omitted or additional ones can be added.

In practice, characterizing $\mathbb{U}^0(x)$ may be difficult, but it is not necessary. After solving the zone tracking problem $\mathbb{P}_N(x)$ to determine the optimal tracking objective value $V_N^0(x)$, a constraint can be added to ensure that the solution to the second optimization achieves the same cost

$$\begin{aligned}\min_{\mathbf{u}\in\mathcal{U}_N} \bar{V}_N(x,\mathbf{u}) = \min_{\mathbf{u}\in\mathcal{U}_N} \quad & \sum_{k=0}^{N-1} c'u + \frac{1}{2}u'Ru + \frac{1}{2}\Delta u'S\Delta u \\ \text{s.t.} \quad & V_N(x,\mathbf{u}) = V_N^0(x)\end{aligned}$$

This approach is formalized in the next section.

7.6.1 Optimization Problem for Zone Tracking

Let $\ell(x)$ be the tracking stage cost satisfying Assumption 7.1 and let $\bar{\ell}(x,u)$ be the economic stage cost. At each time step, the MPC optimization problem solved is

$$\begin{aligned}\bar{V}_N^0(x) = \min_{\mathbf{u}\in\mathcal{U}_N} \bar{V}_N(x,\mathbf{u}) = \sum_{k=0}^{N-1} \bar{\ell}(x(k),u(k)) \\ \text{s.t.} \quad x^+ = f(x,u) \\ x(N) \in \mathcal{A} \\ V_N(x,\mathbf{u}) \leqslant V_N^0(x) + \rho\ell(x)\end{aligned}$$

in which $\rho \in [0, 1)$ is an adjustable tuning parameter that determines how much flexibility the economic controller has in reaching the target set by serving as a slack variable in the constraint. Solving this problem requires knowledge of the optimal tracking cost, which can be obtained using

$$\begin{aligned} V_N^0(x) = \min_{\mathbf{u} \in \mathcal{U}_N} V_N(x, \mathbf{u}) = \sum_{k=0}^{N} \ell(x(k)) \\ \text{s.t.} \quad x^+ = f(x, u) \\ x(N) \in \mathcal{A} \end{aligned}$$

This approach retains the proven stability property and permits additional control over the selection of the optimal policy inside the zone.

7.6.2 Other Approaches

Note that this method requires solving two optimization problems at each MPC step, which may not be desirable. For most practical applications, it may be sufficient to substitute $\bar{\ell}(x, u)$ for $\ell(x, u)$ directly in $\mathbb{P}_N(x)$ provided that the zone tracking term dominates the other terms ($Q \gg c, R, S, \dots$); however, there may not be a theoretical guarantee for stability for this problem if the solution to the resulting optimization problem lies outside of $\mathbb{U}^0(x)$. It is possible that the controller never reaches the target set due to a competing tradeoff in the objective function between the tracking and non-tracking terms.

If all variables are continuous, exact penalties can be used to ensure that the solution does not leave the set $\mathcal{A}$. In this case, the two-step optimization becomes a single optimization problem with an exact penalty on the deviation of the state from the set $\mathcal{A}$. The resulting solution lies in $\mathbb{U}^0(x)$ if the penalty coefficient is chosen to be large enough. With discrete variables, the penalty function may become

discontinuous. However, exact penalties may result in large control actions when a disturbance causes the system to leave the zone, which can be undesirable in practical applications.

7.7 Applications

In order to demonstrate the effectiveness of the zone tracking MPC formulation presented, these concepts are applied to two specific examples: buffer tank level control and commercial building temperature control.

In the following two examples, the tracking term for the stage cost function is given by

$$\ell(x, u) = Q\,|x|_{\mathcal{A}}$$

in which $Q > 0$. The secondary objective, $\bar{\ell}(x, u)$, is modified based on the application.

7.7.1 Buffer Tank Control

In the first example, a buffer tank between upstream and downstream processing units is considered, as shown in Figure 7.5. The goal of the buffer tank is to reject upstream flow disturbances so that they do not propagate downstream. The height of material in the buffer tank is not tracked to a level setpoint, but it should be kept within physical bounds to prevent emptying or overflowing.

The variables for the problem are outlined in Table 7.1. Given an upstream flowrate, F_{in}, the controller manipulates the output flowrate, F_{out}, to maintain the tank level, h, within bounds. The tank level should be maintained between 0.3 m and 0.7 m to provide some cushion so the tank is not in danger of emptying, which can lead to pump cavitation, or overfilling, which can cause backflow to upstream units

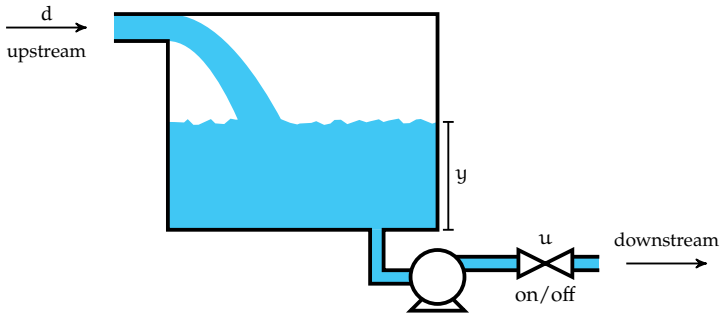

Figure 7.5: Buffer tank process diagram

Table 7.1: Buffer Tank Control Variables

Symbol	Variable	Description	Unit
u	F_{out}	Outlet Flowrate	m^3/s
y	h	Tank Level	m
d	F_{in}	Upstream Flowrate	m^3/s

from pressure buildup. Hence, $\mathcal{A} = [0.3, 0.7]$. The state is assumed to be measured directly in this problem, $y = x$.

Since the goal is to reject upstream flow disturbances by minimizing large variations in the outlet flowrate, the secondary control objective of the zone tracking problem, $\bar{\ell}$, is chosen to penalize the rate of change of the outlet flowrate. For comparison with the zone tracking formulation, a loose setpoint tracking problem is solved. The loose setpoint tracking formulation is given a setpoint of 0.5 m to reach, and it uses a single optimization problem to tradeoff the tracking term with the rate of change penalty term. These two terms are weighted appropriately and tuned based on normal process operating conditions to achieve desirable behavior.

Results. The results from the loose setpoint tracking problem and the zone tracking problem are given in Figure 7.6. The tank is initially empty. At the beginning, the loose setpoint tracking MPC controller is slow to respond when outside of the target set. It takes several hours to drive the tank level to the target set, while the zone

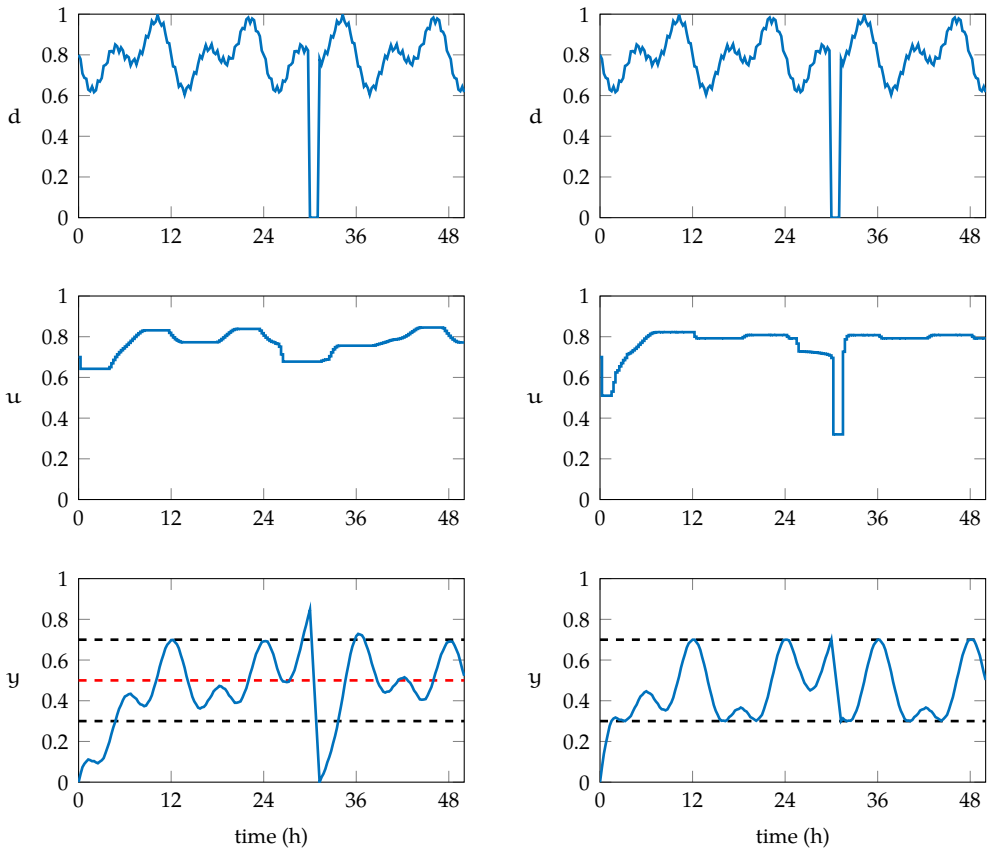

(a) Loose Setpoint Tracking (b) Zone Tracking

Figure 7.6: Comparison of loose setpoint tracking vs zone tracking formulations for buffer tank control

tracking MPC controller gets there quickly. This behavior is expected since the loose setpoint tracking controller is not aggressive when tracking.

Once inside the target set, both controllers perform well in maintaining the level under normal process conditions. The loose setpoint tracking controller was tuned appropriately for these conditions. However, the zone tracking MPC controller does a better job of reducing the upstream flowrate oscillations from propagating downstream since the u profile is flatter.

On the second day, there is a scheduled maintenance of the upstream for one

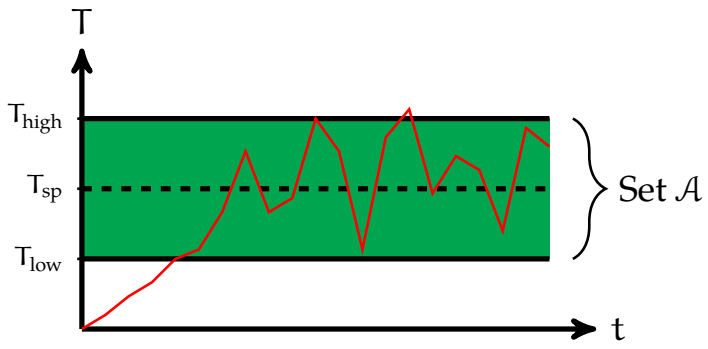

Figure 7.7: Building temperature control

hour. The loose setpoint tracking controller is not able to take enough action to prevent the tank from emptying, however the zone tracking MPC controller is able to maintain the tank level within bounds and then resume nearly constant output flowrate once the upstream comes back online.

Therefore, using loose setpoint tracking can result in undesirable behavior such as emptying the tank under unusual operating conditions for which the controller was not tuned. The zone tracking MPC controller stabilizes the system first, then achieves economic goals. It is able to handle a wide range of operating conditions without the need for constant retuning.

7.7.2 Building Temperature Control

In the second example, the building temperature control problem is addressed. Zone temperatures in building systems do not necessarily need to be at a fixed setpoint. The aim is often to maintain the air temperature between bounds to keep the occupants in the building comfortable, as depicted in Figure 7.7. If the temperature is outside of this zone, the controller should take action to track this "comfort" zone. When the temperature is inside the "comfort" zone, the controller should minimize energy usage, assuming a constant energy cost pricing structure.

Table 7.2: Building Temperature Control Variables

Symbol	Variable	Description	Unit
u	$\dot{Q}_c$	Cooling Delivered from HVAC System	kW
y	T	Zone Temperature	°C
d_1	T_{amb}	Ambient Temperature	°C
d_2	$\dot{Q}_{load}$	Disturbance Load	kW

The variables for the problem are outlined in Table 7.2. Given a forecast of the ambient temperature, T_{amb}, and disturbance load, $\dot{Q}_{load}$, the controller manipulates the cooling delivered from the HVAC system, $\dot{Q}_c$, to regulate the zone temperature, T. The zone temperature should be maintained between 20.5 and 22.5 °C for occupancy comfort, yielding $\mathcal{A} = [20.5, 22.5]$. The state is assumed to be measured directly in this problem, $y = x$.

Since the goal is to minimize energy usage when inside $\mathcal{A}$, the secondary control objective of the zone tracking problem, $\bar{\ell}$, is chosen to be a linear cost on the energy usage of the HVAC system. Similarly, as in the buffer tank problem, a loose setpoint tracking problem is solved for comparison with the zone tracking formulation. The loose setpoint tracking formulation is given a setpoint of 22.4 °C to reach, and it is tuned based on normal building operating conditions to tradeoff the tracking term with the energy usage term. The setpoint given is just inside the comfort zone and maximized in order to reduce energy usage for a more equitable baseline comparison.

Results. The results from the loose setpoint tracking problem and the zone tracking problem are given in Figure 7.8. Both controllers are able to maintain the temperature inside the comfort region under these conditions. While both profiles look similar, the zone tracking MPC controller has an energy usage of 20.8 kWh during this two-day period, while the loose setpoint tracking controller has an energy usage of 29.3

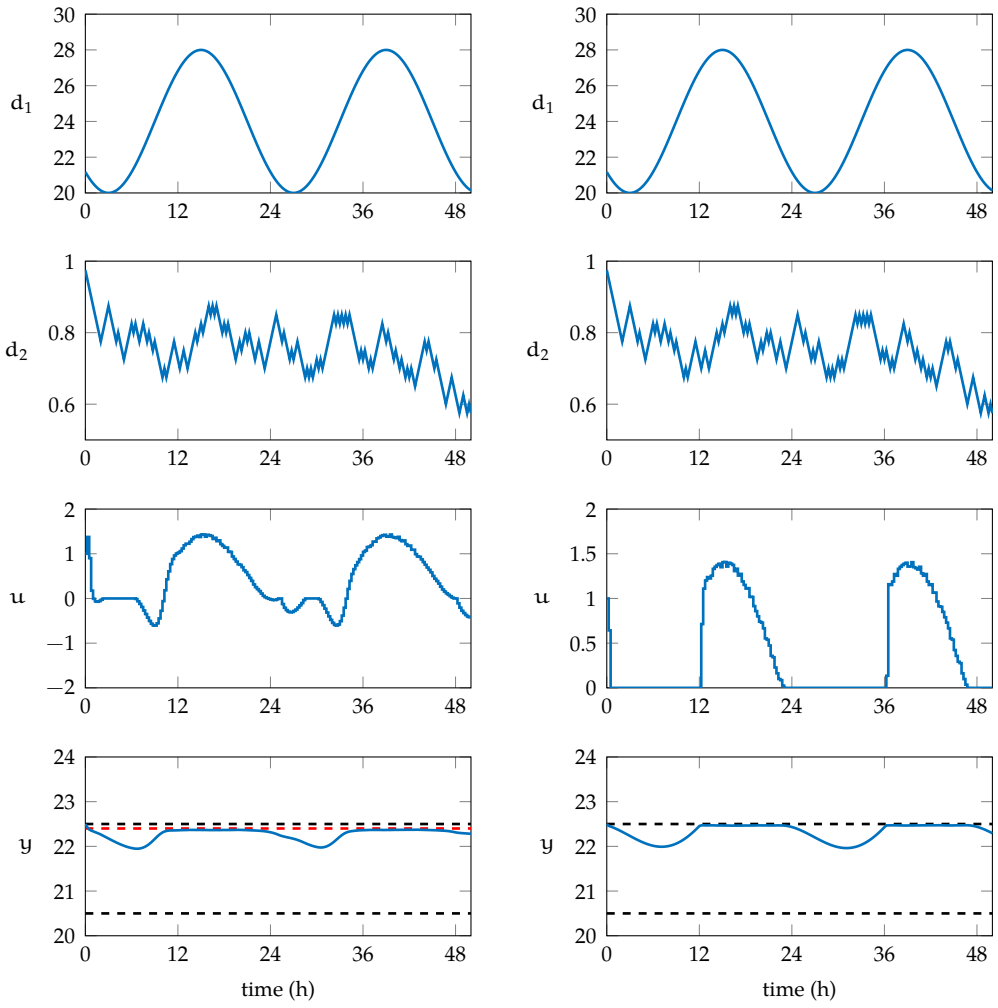

(a) Loose Setpoint Tracking. (29.3 kWh)

(b) Zone Tracking. (20.8 kWh)

Figure 7.8: Comparison of loose setpoint tracking vs zone tracking formulations for building temperature control

kWh. The energy savings of 29% are achieved by the zone tracking controller due to its ability to handle the one side constraint. If a setpoint of 22.5 °C is given to the loose setpoint tracking controller to reduce energy usage, it struggles to respect the upper bound of the comfort region in the present of disturbances. The loose setpoint tracking MPC controller is trading off between tracking the setpoint and minimizing energy usage even when inside the comfort zone. However, it can converge to the

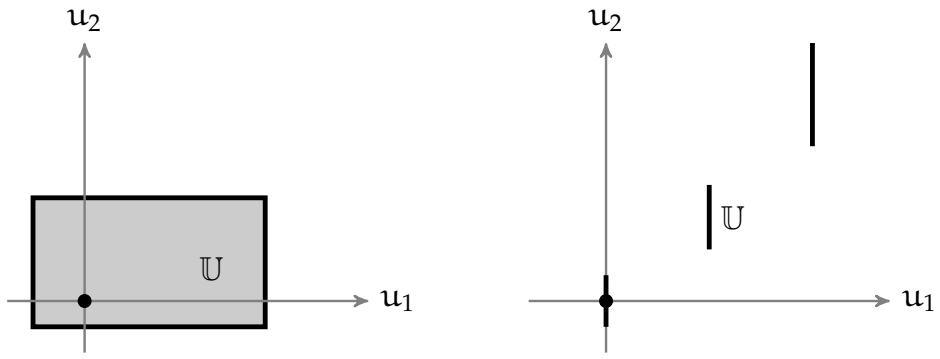

(a) Continuous Actuators (b) Discrete Actuators

Figure 7.9: Constraint sets with continuous and discrete actuators [81]

optimal economic solution of the zone tracking MPC controller in the limit the tracking penalty is zero and exact constraints are enforced on the comfort zone.

7.8 Extension to Discrete Actuators

In the two examples discussed above, the controller makes continuous decisions u. However, not all actuators are continuous. Some control decisions are discrete by nature. For example, on/off valves can take a value of only 0 (off) or 1 (on), not 0.2 or 0.5. Staged HVAC units can deliver only certain rates of cooling as opposed to a full continuum. In this section, this class of systems with discrete actuators is addressed.

When discrete actuators are added to the problem, the primary change occurs in the nature of the input constraint set as shown in Figure 7.9. With continuous actuators, the constraint set often has an interior. However with discrete actuators, an interior is not present, and the value of the control is always at an active constraint. The presence or absence of an interior turns out not to be an issue.

Assumption 7.4 states that the input feasible set $\mathbb{U}$ is compact. $\mathbb{U}$ is not required to have an interior. This key observation, made by Rawlings and Risbeck [81], leads to a natural extension to handle the case of discrete actuators based on the presented

formulation of the zone tracking MPC problem.

Theorem 7.8. *[Folk Theorem]. Any result that holds for standard MPC holds also for MPC with discrete actuators.*

Therefore, the stability theory presented earlier holds even when there are discrete actuators present in the problem [68].

7.9 Applications with Discrete Actuators

The two examples discussed previously in this chapter are revisited, this time with their discrete analogues.

7.9.1 Discrete Buffer Tank Control

In the discrete version of the problem, the control valve at the outlet is replaced by an on/off valve. Physically, this setup represents a system in which the downstream processing unit can be either on or off: $u = \{0, 1\}$. When it is on, it operates at a fixed constant rate. The downstream unit should not be turned on and off too frequently to lengthen the life of the equipment. The objective is similar in that the buffer tank acts to reject the upstream flow disturbances by minimizing the number of nonzero control moves while keeping tank level within bounds. The rest of the problem setup remains identical as previously presented.

For comparison with the zone tracking MPC, *hysteresis control* is used in place of the loose setpoint tracking MPC for the discrete version of the problem. With hysteresis control (also known as *bang-bang control*), the valve is opened when the upper bound is reached and closed when the lower bound is reach. The control algorithm is stated as follows:

1. If $y < y_{\text{low}}$, close valve ($u = 0$).

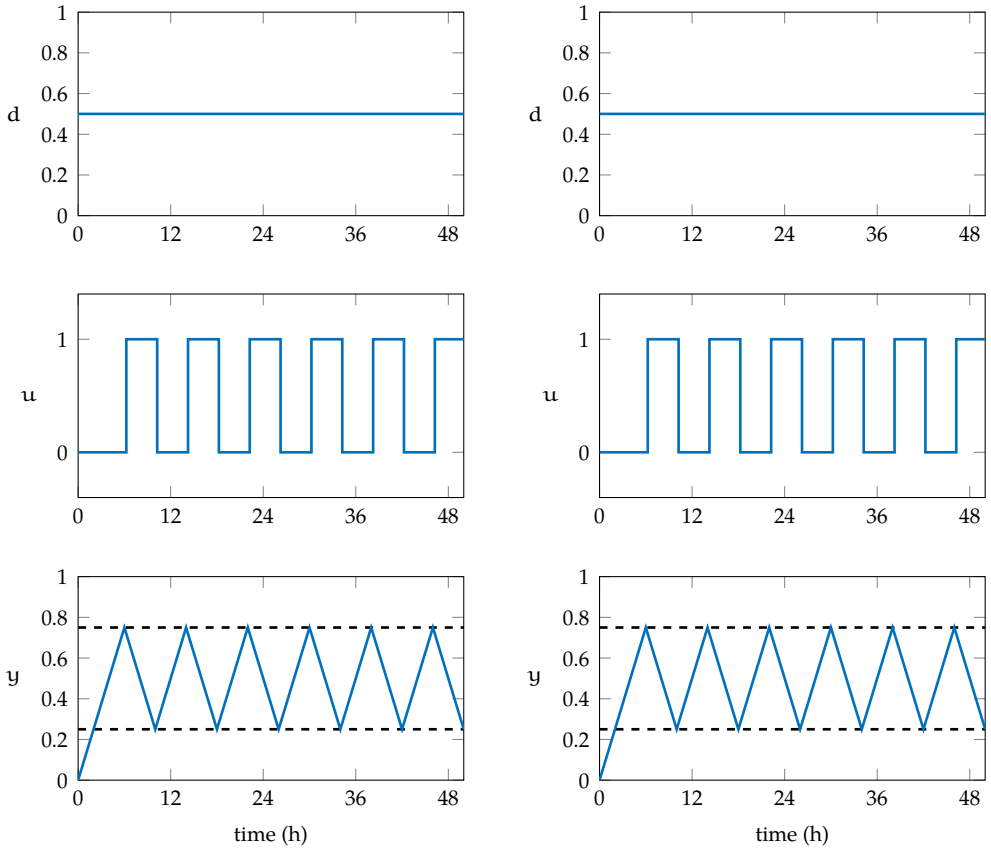

(a) Hysteresis Control (b) Zone Tracking

Figure 7.10: Comparison of hysteresis control vs zone tracking formulations for buffer tank control under ideal conditions

2. If $y > y_{high}$, open valve ($u = 1$).

3. Otherwise, maintain previous valve position ($u = u_{past}$).

This methodology creates a "deadband" for the process variables in place of a setpoint when there are discrete actuators present.

Results. First, the system is simulated under ideal conditions with a constant upstream flowrate and the results are given in Figure 7.10. Under these conditions, both controllers perform identically. The downstream unit is turned on when the

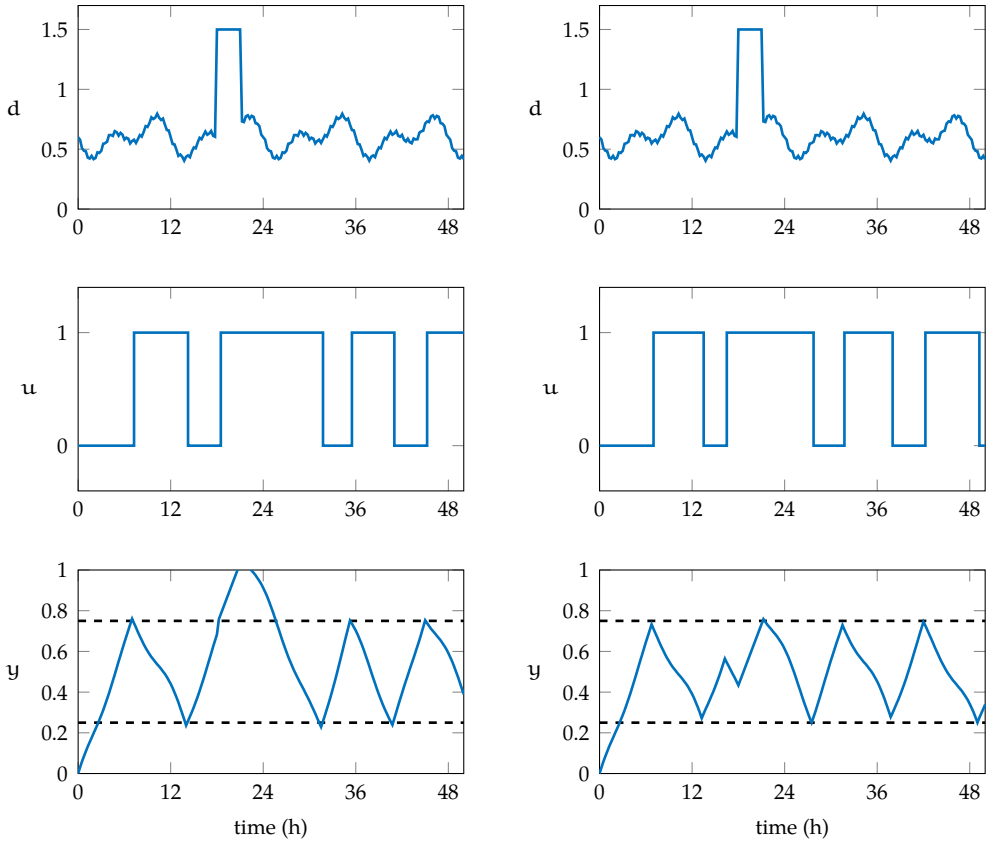

(a) Hysteresis Control (b) Zone Tracking

Figure 7.11: Comparison of hysteresis control vs zone tracking formulations for buffer tank control under nonideal conditions

tank level reaches the lower bound and turned off when it reaches the upper bound. The optimal solution to minimize the number of input moves while maintaining the level within bounds is exactly the hysteresis controller, which is captured by the zone tracking MPC controller.

Next, a more realistic upstream flowrate is applied to the system. The results from this simulation are given in Figure 7.11. Under the normal conditions in the first 12 h, both controllers perform similarly. However, when the upstream unit runs at full capacity for at $t = 18$ for a duration of 3 h, the tank overfills under hysteresis

controller because it does not take action until the upper bound on tank level is reached, by which point, it is too late. In contrast, the zone tracking MPC controller takes early action and runs the downstream unit before this period to drain the tank so it is able to accommodate the excess volume from the upstream unit during this 3 h period.

Therefore, while hysteresis control minimizes number of switches during ideal conditions, it can result in undesirable behavior such as overfilling the tank under unusual operating conditions. However, zone tracking MPC achieves the best of both worlds: it mimics hysteresis control during ideal conditions, while also outperforming hysteresis control over a wide range of operating conditions due to its ability to forecast ahead and optimize.

7.9.2 Discrete Building Temperature Control

In the discrete version of the problem, the continuous HVAC equipment is replaced with two-stage equipment. With two-stage equipment, the HVAC system is able to deliver two rates of cooling depending on how many of its two stages are utilized. The constraint set for the decision v along with the associated cooling rate can be characterized as

$$v \in \{0, 1, 2\} \qquad\qquad u = \dot{Q}_c = v\dot{Q}_{stage}$$

The objectives remain the same: keep temperature between bounds to keep occupants comfortable and minimize energy usage when inside the "comfort" zone. With staged equipment, it is also desirable to reduce the number of switches between the stages whenever possible to lessen the wear and tear on equipment. These terms are considered as secondary objectives in $\bar{\ell}$ when solving the zone tracking MPC problem.

For comparison with the zone tracking MPC, *rule-based control (RBC)* is used since there are two stages in this system requiring integer variables and hysteresis control is designed for binary decisions. The RBC algorithm is defined as follows:

- Operate first stage using hysteresis control:
  1. If $T < T_{low}$, turn off AC ($u = 0$).
  2. If $T > T_{high}$, turn on first stage ($u = 1$).
  3. Otherwise, maintain the previous state ($u = u_{past}$).
- Turn on second stage ($u = 2$) if the following conditions are met:
  - The first stage is on: $u = 1$
  - Runtime of the first stage exceeds threshold: $t_1 > t_{threshold}$
  - Rate of change of temperature is below threshold: $-\frac{dT}{dt} < \dot{T}_{threshold}$

This approach serves as an extension of the hysteresis control and is one commonly used method in building systems.

Results. The results from this study are provided in Figure 7.12. For this two-day period, the RBC method achieves energy usage of 85.5 kWh, and the zone tracking MPC controller uses 75.0 kWh of energy. RBC relies on comfort zone violations to trigger a switch in the state of the AC system. As with the hysteresis controller, this reactive approach can lead to bounds being violated; the RBC is not able to maintain the temperature within comfort bounds when a large disturbance enters the system at $t = 36$ since it does not forecast into the future. It uses more energy and does not minimize switching. By contrast, the zone tracking MPC controller maintains the temperature in the comfort zone with fewer switches and achieves energy savings of 12% compared to the rule-based controller.

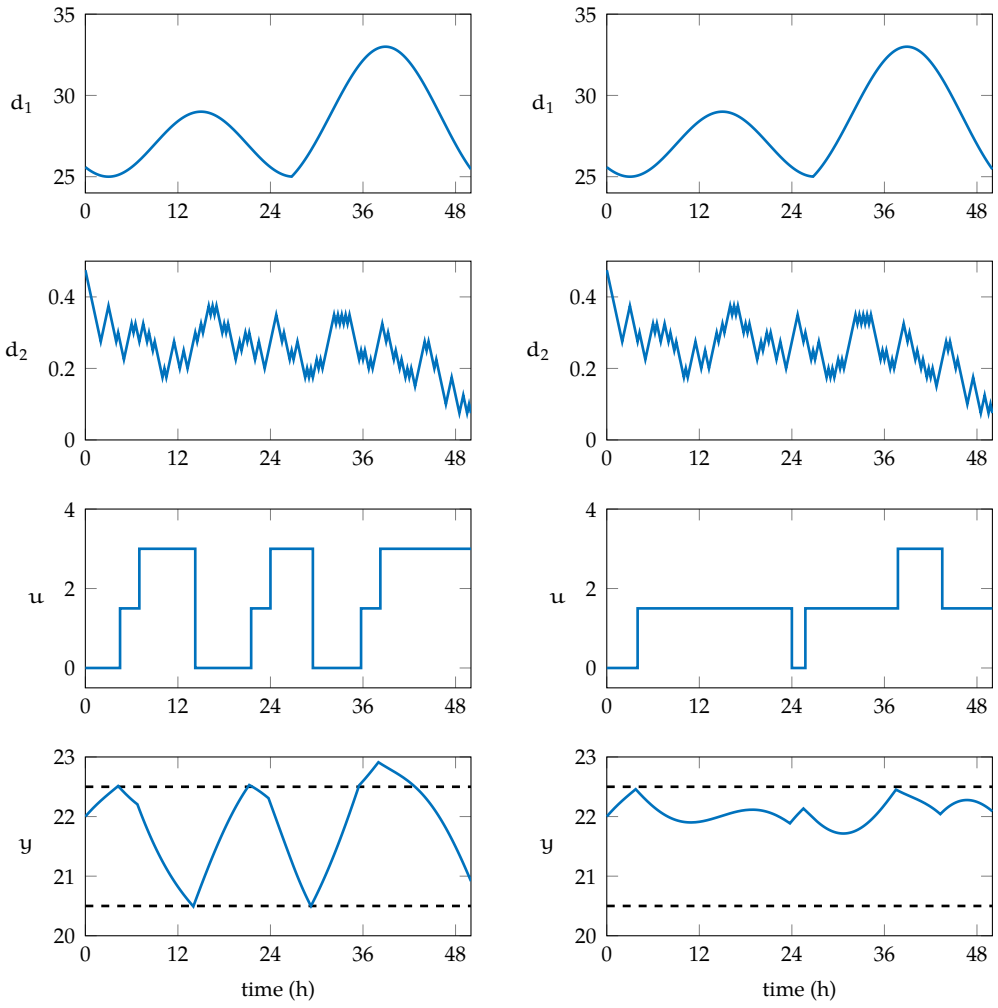

(a) Rule-Based Control. (85.5 kWh)

(b) Zone Tracking. (75.0 kWh)

Figure 7.12: Comparison of rule-based control vs zone tracking formulations for building temperature control

Furthermore, large-scale commercial building applications contain hundreds of zones and many pieces of HVAC equipment as shown in Figure 2.1. It is difficult to develop rule-based heuristics for these applications when there are thousands of continuous and discrete decisions to make. The zone tracking MPC formulation generalizes easily to handle any size system since it can be formulated as an optimization problem.

7.10 Summary

A formulation for tracking zones using MPC has been presented. In zone tracking, the tracking stage cost is zero if the state is in the target set and positive otherwise. Under the stated assumptions, the target set is shown to be asymptotically stable for the closed-loop system even in the presence of discrete actuators. The proposed formulation also addresses secondary objectives. Since the solution of the optimization problem is often nonunique when tracking zones, a general framework is proposed to achieve secondary control objectives such as minimizing economic costs, input value, or input movement. After solving the zone tracking problem, a second optimization problem can be solved to determine a unique input trajectory from the stabilizing set that optimizes secondary control objectives. Furthermore, the theory extends to handle the case of discrete actuators.

The effectiveness of this formulation is demonstrated on two example problems: buffer tank control and building temperature control. The zone tracking MPC formulation has many advantages:

- Reduces to the hysteresis controller in the ideal case
- Outperforms the loose setpoint tracking MPC and hysteresis controllers over a wider range of operations (buffer tank example)
- Optimizes economics in real-time (building temperature example)
- Generalizes to any number of discrete decisions

The zone tracking MPC controller is shown to track target zones first then optimize secondary objectives afterwards in both cases.

8 CONCLUSIONS

Wide-scale implementation is key to realizing the tremendous energy cost benefits of optimizing resource usage in building HVAC applications. Early projects, such as SESI, have demonstrated the economic benefits to consumers. The optimization of energy storage resources through load shifting and peak demand management creates a flatter, more constant load for power plants which allows for more efficient usage and conversion of primary energy at the grid scale. Hence, both the supplier and the consumer benefits in this arrangement.

An optimization-based system such as MPC has been shown to be a promising solution to manage these integrated resources operating on multiple time-scales. While significant progress has been achieved to date to identify major players such as university campuses and industrial complexes, the pre-existing MPC technology is not able to address the full airside and waterside problem for these types of applications. Several obstacles remain in the design and implementation of the control architecture. In this work, several steps have been taken towards overcoming these obstacles. The major contributions of this work are summarized in this chapter, and future directions for this research area are explored.

8.1 Summary of Contributions

Hierarchical Decomposition. In Chapter 3, we propose a novel hierarchical decomposition for the large-scale commercial HVAC problem defined in Chapter 2. By decomposing the overall combined problem into smaller subproblems, the computational complexity is decreased. Aggregate system models are used to facilitate this hierarchical decomposition. A key design feature of this control system is that iterations are not required, which reduces the communication burden and ensures

that solution times are not increased by the presence of delays between distributed controllers. Peak demand charge is also handled by the high-level coordinating layer. The proposed control system can achieve 5–40% energy cost savings depending on the amount of energy storage available in the application and the market incentives for load shifting.

Implementation Strategies. In Chapter 4, the development of building models, a key step in the implementation of such control systems, is addressed. We recommend a grey-box system identification procedure over the black-box approach to reduce the amount of data needed, which in turn decreases the costs of installation and commissioning these procedures. Increasing the complexity of the model may not necessarily increase the savings.

As central plant equipment becomes more efficient, airside power consumption is starting to become a more significant portion of the total energy costs. The load due to airside power consumption cannot be shifted using active TES. Hence passive TES must be used to shift part of the load. The more efficient method is to use electrical storage for this purpose.

Battery Storage. In Chapter 5, we extend the proposed framework to include electrical storage for embedded battery applications. The storage utilization decision is distributed amongst airside controllers with a high-level coordinator managing the peak demand charge. With the cost of manufacturing batteries decreasing, these HVAC units have the potential to become financially lucrative. Battery storage also opens the door to participation in electricity market revenue generation programs.

Case Study. In Chapter 6, we developed a case study loosely based on the Stanford University central HVAC plant. The campus includes both an airside and waterside

system. The airside consists of 500 zones and the waterside consists of several chillers, pumps, and cooling towers. We demonstrated that the proposed hierarchical decomposition can be solved in real-time, even for this large system. Some of the equipment models and data were obtained directly from our industrial collaborator, Johnson Controls. Hence, the valuable models and data for this study are made publicly available online for other researchers in the HVAC community to propose and evaluate other control architectures. The files can be downloaded from the following website: `https://hvacstudy.github.io/`.

Zone Tracking MPC. In Chapter 7, a novel procedure for tracking zone regions is outlined. The approach relies on stabilizing the zone region first, then achieving secondary economic objectives. Nominal stability theory for this control methodology is proven. Discrete actuators are addressed as well. The advantages of the proposed framework over conventional setpoint tracking is demonstrated via simulation examples. In the ideal case, the zone tracking MPC reduces to the hysteresis controller. However, it generalizes and performs well in handle nonideal cases too. This zone tracking MPC approach can be used as the tracking regulatory layer, replacing PI controllers in building temperature regulation.

8.2 Future Directions

There are several directions for future research in the area of commercial HVAC control. Some of them include disturbance forecasting, state estimation, the replacement of the regulatory layer with MPC and handling of discrete decisions in the airside system.

Disturbance Forecasting and State Estimation. The areas of disturbance forecasting and state estimation have not received much attention in the HVAC community [69]. The generation of load forecasts for use in MPC as feedforward variables is vital in determining policies to utilize energy storage resources optimally. Historical data must be used in conjunction with weather forecasting to make these forecasts. Auxiliary information such as daily working hours or meeting schedules for intermittently occupied conference room spaces can also be used to determine the effect of occupancy on heat generation. There is an opportunity in this area to develop methods which can use these input data to generate load forecasts for each zone.

State estimation is another interesting area of research. Models and predictions are not perfect, especially in buildings. When measurements are taken of the system, methods of updating the disturbance estimate are required. More work is needed to determine whether previously proposed state estimation techniques and disturbance models can perform well with the large slow periodic disturbances present in building temperature regulation. It is also not clear what the optimal procedure may be for using this single instant measurement information to correct disturbance forecast *trajectories* for use in the MPC problem.

Replacement of Regulatory Layer. MPC has been treated in this work as a supervisory controller that sends setpoints to regulatory controllers. This framework has wide applicability for existing buildings where the building automation infrastructure is difficult to displace. However, with the construction of new commercial buildings, there is an opportunity to escape this paradigm. Rather than MPC deciding the setpoints for regulatory controllers in order to manipulate their response behavior for decreasing energy costs, MPC can make those regulatory decisions directly. This approach also removes the need to model regulatory controllers in the MPC layer, a major source of nonlinearity. Regulatory controllers in buildings can

have adaptive seasonal tuning and control logic which can be difficult to represent with linear models suitable for optimization. One candidate to replace these regulatory controllers may be the zone tracking MPC controller presented in this work, which is able to track zone regions while optimizing secondary economic objectives.

There are challenges associated with replacing these controllers. MPC making lower level decisions is required to run on a faster time scale than every 15 minutes. The solution times of the optimization problems would need to be less than the regulatory control execution step, on the order of seconds. Exploration of this idea can elucidate whether the added problem complexity is warranted by the savings improvement.

Discrete Decisions in Airside System. The discrete decisions are primarily treated for the waterside system only, e.g., when scheduling equipment operation. However, there may be some applications where discrete decisions appear in the airside problem as well. One example is the cooling or heating mode decision in VRF systems. If the existing regulatory layer is replaced with MPC as suggest above, then MPC may have to make discrete on/off decisions for staged HVAC equipment on the airside. The effect of this change on the choice of decomposition and its scalability needs further exploration.

Commercial HVAC systems in buildings have been a rich application involving several disciplines, including traditional process control for the airside system, planning and scheduling for waterside equipment operation, and stochastic nature of variables, such as electricity pricing and disturbance loads. As new research areas arise, a multidisciplinary approach is going to be required to achieve the desired objectives.

A APPENDIX I: FORMULATION

In this appendix, we present additional details concerning the formulation of the MPC optimization problems that are omitted from the main text.

The proposed decomposition approach from Chapter 3 is shown in Figure A.1. The high-level problem uses aggregate models of the airside subsystems. First, a procedure for aggregation of airside zone models into subsystem (e.g., building) models is given. Then, the formulation of the high-level and low-level airside MPC problems as optimization problems in standard form is shown. Certain constraints, such as comfort bounds, are implemented as *soft constraints* by adding slack variables and penalizing their violations in the objective function to ensure the optimizations are always feasible.

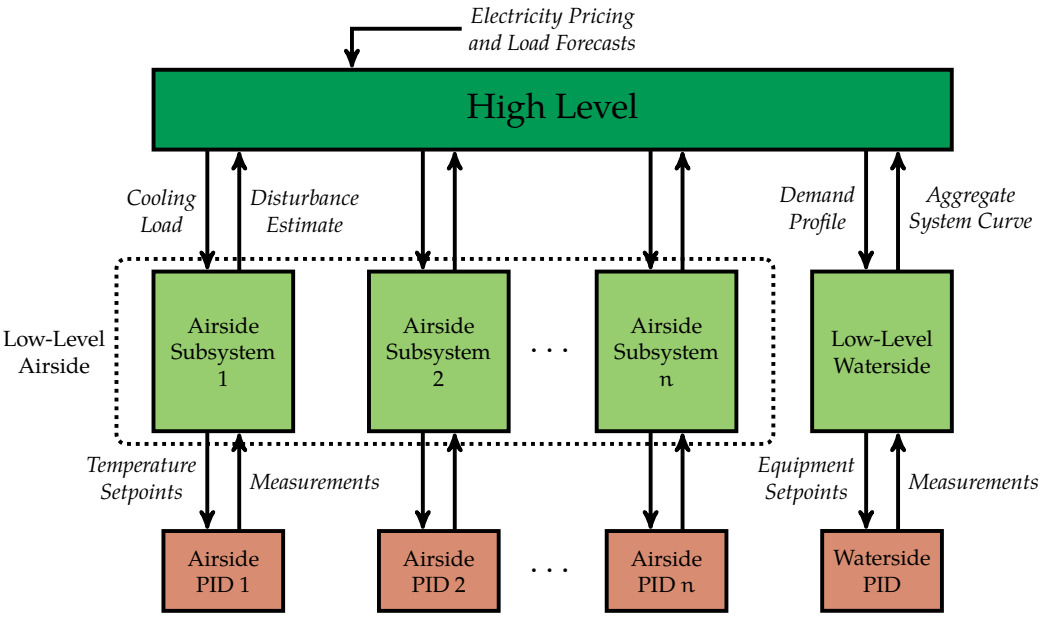

Figure A.1: Hierarchical decomposition with distributed low-level airside subsystems

A.1 Airside Model Aggregation

Once model parameters have been identified for the low-level airside system, they can be aggregated to form a model for the high-level problem. The energy balance for the low-level airside system is given by

$$C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{other,i}$$

Taking the sum over all zones i in a building gives

$$\sum_i C_i \frac{dT_i}{dt} = -\sum_i H_i(T_i - T_a) - \sum_i \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \sum_i \dot{Q}_{c,i} + \sum_i \dot{Q}_{other,i}$$

This equation can be simplified term-by-term as follows.

1. Since the temperatures of the zones do not vary much (width of the comfort zone is only 2°C), the heat capacity is assumed to be constant. Assuming that the zones are closed systems with no mass entering or leaving, C_i is constant. With constant C_i, the accumulation term can be expressed in terms of the building heat capacitance and average building temperature, T_b, which is the steady-state equilibrium temperature of the isolated system.

$$\begin{aligned} \sum_i C_i \frac{dT_i}{dt} &= \sum_i \frac{d(C_i T_i)}{dt} = \frac{C_b}{C_b} \sum_i \frac{d(C_i T_i)}{dt} \\ &= C_b \frac{d\left(\sum_i \frac{C_i T_i}{C_b}\right)}{dt} = C_b \frac{dT_b}{dt} \end{aligned}$$

2. We would like to have an analagous heat transfer term with the ambient in

terms of the average building temperature

$$-\sum_i H_i(T_i - T_a) = -H_b(T_b - T_a)$$

Substituting T_b and solving for H_b gives

$$\sum_i H_i(T_i - T_a) = H_b(T_b - T_a)$$
$$H_b = \frac{\sum_i H_i(T_i - T_a)}{(T_b - T_a)}$$
$$H_b = \frac{\sum_i H_i(T_i - T_a)}{\left(\sum_i \frac{C_i T_i}{C_b} - T_a\right)}$$

The aggregate scaled heat transfer coefficient, H_b, is a function of the zone temperatures, T_i and ambient temperature, T_a in addition to the zone model parameters. Hence, the resulting aggregate state-space model is time-varying.

To obtain a time-invariant model for the controller, the zone temperatures T_i are assumed to be approximately equal. This approximation is valid for HVAC systems since the comfort regions for zone temperature are relatively narrow. With all zones having the same temperature, we have that $T_i \approx T_b$ for all zones i, which gives

$$H_b = \frac{\sum_i H_i(T_i - T_a)}{(T_b - T_a)} \approx \frac{\sum_i H_i(T_b - T_a)}{(T_b - T_a)}$$
$$H_b \approx \sum_i H_i$$

This approximate aggregate heat transfer coefficient is used in the controller.

3. The average building temperature does not change with heat transfer between zones since that occurs within the system boundary (i.e., the coupling terms

cancel each other).

$$\sum_{i} \sum_{j \neq i} \beta_{ij}(T_i - T_j) = 0$$

This fact is best illustrated with an example. With $\beta_{ij} = \beta_{ji}$, consider a four-zone building:

$$\begin{aligned}\sum_{i} \sum_{j \neq i} \beta_{ij}(T_i - T_j) = {} & \beta_{12}(T_1 - T_2) + \beta_{13}(T_1 - T_3) + \beta_{14}(T_1 - T_4) \\ & + \beta_{21}(T_2 - T_1) + \beta_{23}(T_2 - T_3) + \beta_{24}(T_2 - T_4) \\ & + \beta_{31}(T_3 - T_1) + \beta_{32}(T_3 - T_2) + \beta_{34}(T_3 - T_4) \\ & + \beta_{41}(T_4 - T_1) + \beta_{42}(T_4 - T_2) + \beta_{43}(T_4 - T_3) \\ = {} & 0\end{aligned}$$

4. Cooling duty is an input and other loads are disturbances, so we simply sum them.

$$\sum_{i} \dot{Q}_{c,i} = \dot{Q}_{c,b} \qquad \sum_{i} \dot{Q}_{other,i} = \dot{Q}_{other,b}$$

After simplifications, the energy balance for the high-level problem becomes

$$C_b \frac{dT_b}{dt} = -H_b(T_b - T_a) - \dot{Q}_c + \dot{Q}_{other}$$

This process can be repeated for other buildings.

To summarize, once C_i and H_i have been identified using grey-box identification on the low-level airside system, the aggregate building model parameters are

computed using

$$C_b = \sum_i C_i \qquad H_b = \sum_i H_i$$

with the relationships

$$T_b = \sum_i \frac{C_i T_i}{C_b} \qquad \dot{Q}_{c,b} = \sum_i \dot{Q}_{c,i} \qquad \dot{Q}_{other,b} = \sum_i \dot{Q}_{other,i}$$

A.2 High Level Problem

The data for the high-level problem includes electricity pricing, weather forecast, building use model, and disturbance estimate from low-level problem. Its objective is to minimize cost over the prediction horizon. The decision variables are average building temperatures, storage utilization, and thermal loads for waterside system. The models required include the average building temperature to heat duty model, active thermal energy storage model, and aggregate waterside system model. The constraints include comfort bounds on the air temperature; minimum, maximum, and rate of change of the heat duty; and minimum, maximum, and rate of change of the electrical power.

A.2.1 Optimization Problem

With the inclusion of the airside power consumption model, the high-level MPC problem is formulated as

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u},\dot{W}_{\text{peak}}} \quad & \sum_{k=0}^{N-1} c_k \left(\eta_{\text{HVAC}} \dot{Q}_{\text{HVAC},k} + \eta_{\text{air}} \sum_b \dot{Q}_{c,b,k} \right) \Delta + c_{\text{peak}} \dot{W}_{\text{peak}} \\
\text{s.t.} \quad & C_b \frac{dT_b}{dt} = -H_b (T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{\text{other},b} \\
& \frac{ds}{dt} = -\sigma s + \dot{Q}_{\text{storage}} \\
& \dot{Q}_{\text{HVAC},k} = \sum_b \dot{Q}_{c,b,k} + \dot{Q}_{\text{storage},k} \\
& \left(\eta_{\text{HVAC}} \dot{Q}_{\text{HVAC},k} + \eta_{\text{air}} \sum_b \dot{Q}_{c,b,k} \right) \leqslant \dot{W}_{\text{peak}} \\
& \dot{W}_{\text{peak,past}} \leqslant \dot{W}_{\text{peak}}, 0 \leqslant \dot{Q}_{\text{HVAC},k} \leqslant \dot{Q}_{\text{HVAC,max}} \\
& 0 \leqslant s_k \leqslant s_{\max} \\
& T_{\min} \leqslant T_b \leqslant T_{\max}
\end{aligned}
$$

The airside operation determines η_{air} and the waterside operation determines η_{HVAC}, both of which are treated as constants in the High Level problem. The demand change constraint has to be modified from the simplified implementation used in the previous formulation.

A.2.2 Formulation

The high-level problem is formulated as a linear program of the form:

$$
\min_z c^T z \qquad \text{s.t.} \quad Ez = e, \quad Fz \leqslant f, \quad l \leqslant z \leqslant u
$$

States	$x = \begin{bmatrix} T_b & s \end{bmatrix}^T$	$n = n_b + 1$
Inputs	$u = \begin{bmatrix} \dot{Q}_{c,b} & \dot{Q}_{storage} \end{bmatrix}^T$	$m = n_b + 1$
Measurements	$y = \begin{bmatrix} T_b & s \end{bmatrix}^T$	$p = n_b + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,b} \end{bmatrix}^T$	$n_d = 1 + n_b$

Table A.1: High-Level Problem Variables

The subsequent sections outline the procedure for obtaining a linear program of this form.

Model

As shown in Table A.1, the states are the building temperatures and storage tank level. The inputs are the cooling duties to each building and the amount of storage charged or discharged. All states are assumed to be measured perfectly. The disturbances are the ambient temperature and external loads placed on each building.

The energy balance ($C_b \frac{dT_b}{dt} = -H_b(T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{other,b}$) can be expressed

as a continuous-time state equation for the first n_b states and first n_b inputs.

$$\frac{d}{dt}\begin{bmatrix} T_1 \\ T_2 \\ \vdots \\ T_{n_b} \end{bmatrix} = \underbrace{\begin{bmatrix} \frac{-H_1}{C_1} & 0 & \cdots & 0 \\ 0 & \frac{-H_2}{C_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{-H_{n_b}}{C_{n_b}} \end{bmatrix}}_{A_c} \begin{bmatrix} T_1 \\ T_2 \\ \vdots \\ T_{n_b} \end{bmatrix}$$

$$+ \underbrace{\begin{bmatrix} \frac{-1}{C_1} & 0 & \cdots & 0 \\ 0 & \frac{-1}{C_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{-1}{C_{n_b}} \end{bmatrix}}_{B_c} \begin{bmatrix} \dot{Q}_{c,1} \\ \dot{Q}_{c,2} \\ \vdots \\ \dot{Q}_{c,n_b} \end{bmatrix}$$

$$+ \underbrace{\begin{bmatrix} \frac{H_1}{C_1} & \frac{1}{C_1} & 0 & \cdots & 0 \\ \frac{H_2}{C_2} & 0 & \frac{1}{C_2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{H_{n_b}}{C_{n_b}} & 0 & 0 & \cdots & \frac{1}{C_{n_b}} \end{bmatrix}}_{\Gamma_c} \begin{bmatrix} T_a \\ \dot{Q}_{other,1} \\ \dot{Q}_{other,2} \\ \vdots \\ \dot{Q}_{other,n_b} \end{bmatrix}$$

Discretize the continuous-time model above using the controller interval, Δ, to get the discrete version of the above matrices: $\tilde{A}$, $\tilde{B}$, and $\tilde{\Gamma}$. Then, we can add the storage model ($s_{k+1} = \eta_{storage} s_k + \dot{Q}_{storage,k}\Delta$) to get the final discrete-time state space model:

$$x^+ = Ax + Bu + B_d d$$

$$y = Cx + Du$$

where

$$A = \begin{bmatrix} \tilde{A} & 0 \\ 0 & \eta_{\text{storage}} \end{bmatrix}$$

$$B = \begin{bmatrix} \tilde{B} & 0 \\ 0 & \Delta \end{bmatrix}$$

$$B_d = \begin{bmatrix} \tilde{\Gamma} \\ 0 \end{bmatrix}$$

$$C = I \qquad D = 0$$

Variables

The variables in the linear program are

$$z = \begin{bmatrix} u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \end{bmatrix} \Big\} N(m+n)$$
$$\begin{bmatrix} \dot{W}_{\text{peak}} \end{bmatrix} \Big\} 1$$
$$\begin{bmatrix} p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \end{bmatrix} \Big\} Nm$$
$$\begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_N \end{bmatrix} \Big\} Nn_b$$

The total number of variables is $N(2m + n + n_b) + 1$.

Objective

The linear objective function is

$$
c = \begin{array}{l} \\ u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \\ \dot{W}_{\text{peak}} \\ p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \\ p_1 \\ p_2 \\ \vdots \\ p_N \end{array}
\begin{bmatrix} \\ c_0 \Delta w \\ 0 \\ c_1 \Delta w \\ 0 \\ \vdots \\ c_{N-1} \Delta w \\ 0 \\ c_{\text{peak}} \\ \alpha \\ \alpha \\ \alpha \\ \vdots \\ \alpha \\ \beta \\ \beta \\ \vdots \\ \beta \end{bmatrix}
$$

The duty to power vector, w, is given by

$$w = \eta_{\text{HVAC}} 1_m + \eta_{\text{air}} \begin{bmatrix} 1_{m-1} \\ 0 \end{bmatrix}$$

where 1_m is a column vector of ones with length m. Note that these definitions give

$$w^T u = \eta_{\text{HVAC}} \dot{Q}_{\text{HVAC}} + \eta_{\text{air}} \sum_b \dot{Q}_{c,b,k}$$

Constraints

The state equation ($x^+ = Ax + Bu + B_d d$) is represented as $Ez = e$ where

$$E = \begin{array}{c} \begin{array}{cccccccccccc} u(0) & x(1) & u(1) & x(2) & u(2) & x(3) & \cdots & u(N-1) & x(N) & \dot{W}_{\text{peak}} & \cdots \end{array} \\ \begin{bmatrix} -B & I & & & & & & & & & \\ & -A & -B & I & & & & & & & \\ & & & -A & -B & I & & & & & \\ & & & & & & \ddots & & & & \\ & & & & & & -A & -B & I & & \end{bmatrix} \end{array}$$

$$e = \begin{bmatrix} B_d d(0) \\ B_d d(1) \\ B_d d(2) \\ \vdots \\ B_d d(N-1) \end{bmatrix} + \left.\begin{bmatrix} Ax(0) \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}\right\} Nn$$

The demand charge constraint ($\eta_{\text{HVAC}}\dot{Q}_{\text{HVAC},k} + \eta_{\text{air}} \sum_b \dot{Q}_{c,b,k} \leqslant \dot{W}_{\text{peak}}$) is represented as $F_1 z \leqslant f_1$ where

$$F_1 = \begin{array}{c} \begin{matrix} u(0) & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & x(N) & \dot{W}_{\text{peak}} & \cdots \end{matrix} \\ \begin{bmatrix} w^T & & & & & & & & -1 & \\ & & w^T & & & & & & -1 & \\ & & & & w^T & & & & -1 & \\ & & & & & \ddots & & & \vdots & \\ & & & & & & w^T & & -1 & \end{bmatrix} \end{array}$$

$$f_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \Bigg\} N$$

and 1_m^T is a row vector of ones with length m.

The nonnegativity HVAC constraint ($0 \leqslant \dot{Q}_{\text{HVAC},k}$) is represented as $F_2 z \leqslant f_2$ where

$$F_2 = \begin{array}{c} \begin{matrix} u(0) & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & x(N) & \dot{W}_{\text{peak}} & \cdots \end{matrix} \\ \begin{bmatrix} -1_m^T & & & & & & & & & \\ & & -1_m^T & & & & & & & \\ & & & & -1_m^T & & & & & \\ & & & & & \ddots & & & & \\ & & & & & & -1_m^T & & & \end{bmatrix} \end{array}$$

$$f_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \Bigg\} N$$

The HVAC capacity constraint ($\dot{Q}_{HVAC,k} \leqslant \dot{Q}_{HVAC,max}$) is represented as $F_3 z \leqslant f_3$ where

$$F_3 = \begin{array}{c} \begin{array}{cccccccccc} u(0) & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & x(N) & \dot{W}_{peak} & \cdots \end{array} \\ \begin{bmatrix} 1_m^T & & & & & & & & & \\ & & 1_m^T & & & & & & & \\ & & & & 1_m^T & & & & & \\ & & & & & \ddots & & & & \\ & & & & & & & 1_m^T & & \end{bmatrix} \end{array}$$

$$f_3 = \begin{bmatrix} \dot{Q}_{HVAC,max} \\ \dot{Q}_{HVAC,max} \\ \dot{Q}_{HVAC,max} \\ \vdots \\ \dot{Q}_{HVAC,max} \end{bmatrix} \Bigg\} N$$

and 1_m^T is a row vector of ones with length m.

The rate of change constraint is represented as $F_4 z \leqslant f_4$ where

$$
F_4 = \begin{array}{c}
\begin{array}{ccccccccccccccc}
u(0) & & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & \cdots & p_{0,-1} & p_{1,0} & p_{2,1} & \cdots & p_{N-1,N-2} & \cdots
\end{array} \\
\begin{bmatrix}
I & & & & & & & & & -I & & & & & \\
-I & & & & & & & & & -I & & & & & \\
-I & & I & & & & & & & & -I & & & & \\
I & & -I & & & & & & & & -I & & & & \\
 & & -I & & I & & & & & & & -I & & & \\
 & & I & & -I & & & & & & & -I & & & \\
 & & & & & \ddots & & & & & & & \ddots & & \\
 & & & & -I & & I & & & & & & & -I & \\
 & & & & I & & -I & & & & & & & -I &
\end{bmatrix}
\end{array}
$$

$$
f_4 = \left.\begin{bmatrix} u(-1) \\ -u(-1) \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}\right\} 2Nm
$$

The comfort zone constraint ($T_{min} \leqslant T_i \leqslant T_{max}$) is represented as $F_5 z \leqslant f_5$ where

$$
F_5 = \begin{array}{c}
\begin{array}{ccccccccccccc}
u(0) & x(1) & u(1) & x(2) & \cdots & u(N-1) & x(N) & \cdots & p_1 & p_2 & \cdots & p_N
\end{array} \\
\begin{bmatrix}
 & -G & & & & & & & -I & & & \\
 & G & & & & & & & -I & & & \\
 & & & -G & & & & & & -I & & \\
 & & & G & & & & & & -I & & \\
 & & & & \ddots & & & & & & \ddots & \\
 & & & & & & -G & & & & & -I \\
 & & & & & & G & & & & & -I
\end{bmatrix}
\end{array}
$$

$$f_5 = \left.\begin{bmatrix} -T_{min} \\ T_{max} \\ -T_{min} \\ T_{max} \\ \vdots \\ -T_{min} \\ T_{max} \end{bmatrix}\right\} 2Nn_b$$

and where $G = [I \quad 0]$, T_{min} is a column vector of length z with the lower end of the temperature comfort zone, and T_{max} is a column vector of length z with the upper end of the temperature comfort zone.

The bounds on the variables are

$$l = \begin{bmatrix} u_{min} \\ x_{min} \\ u_{min} \\ x_{min} \\ \vdots \\ u_{min} \\ x_{min} \\ \dot{W}_{peak,past} \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \leqslant \begin{bmatrix} u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \\ \dot{W}_{peak} \\ p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \\ p_1 \\ p_2 \\ \vdots \\ p_N \end{bmatrix} \leqslant \begin{bmatrix} u_{max} \\ x_{max} \\ u_{max} \\ x_{max} \\ \vdots \\ u_{max} \\ x_{max} \\ \infty \\ \infty \\ \infty \\ \infty \\ \vdots \\ \infty \\ \infty \\ \infty \\ \vdots \\ \infty \end{bmatrix} = u$$

where

$$u_{min} = \begin{bmatrix} \dot{Q}_{c,1,low} \\ \dot{Q}_{c,2,low} \\ \vdots \\ \dot{Q}_{c,n_b,low} \\ \dot{Q}_{storage,low} \end{bmatrix} \leqslant \begin{bmatrix} \dot{Q}_{c,1} \\ \dot{Q}_{c,2} \\ \vdots \\ \dot{Q}_{c,n_b} \\ \dot{Q}_{storage} \end{bmatrix} \leqslant \begin{bmatrix} \dot{Q}_{c,1,high} \\ \dot{Q}_{c,2,high} \\ \vdots \\ \dot{Q}_{c,n_b,high} \\ \dot{Q}_{storage,high} \end{bmatrix} = u_{max}$$

$$x_{min} = \begin{bmatrix} -\infty \\ -\infty \\ \vdots \\ -\infty \\ 0 \end{bmatrix} \leqslant \begin{bmatrix} T_1 \\ T_2 \\ \vdots \\ T_{n_b} \\ s \end{bmatrix} \leqslant \begin{bmatrix} \infty \\ \infty \\ \vdots \\ \infty \\ s_{max} \end{bmatrix} = x_{max}$$

By defining

$$F = \begin{bmatrix} F_1 \\ F_2 \\ F_3 \\ F_4 \\ F_5 \end{bmatrix} \qquad f = \begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \end{bmatrix}$$

these constraints can be represented in the form

$$Ez = e$$

$$Fz \leqslant f$$

$$l \leqslant z \leqslant u$$

A.3 Low-Level Airside Problem

The data for the low-level airside problem includes the heat duty over the prediction horizon from the high-level problem and temperature measurements from PID

controllers. Its objective is to minimize energy usage over the prediction horizon. The decision variables are zone temperature setpoints. The zone air temperature and temperature setpoint to heat duty model is required, which includes the model of the PID controllers. The constraints include comfort bounds on the zone air temperatures and bounds on the heat duty, which are obtained from the high-level problem.

A.3.1 Optimization Problem

Mathematically, the low-level airside MPC problem is formulated as

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u}} \quad & Q_{\text{total}} \\
\text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \\
& T_{\min} \leqslant T_i \leqslant T_{\max} \\
& \frac{dQ_{\text{tot}}}{dt} = \sum_i \dot{Q}_{c,i} \\
& \dot{Q}_{c,i} = f(T_i, T_{\text{sp},i}) \\
& \sum_i \dot{Q}_{c,i} \leqslant \dot{Q}_{\text{HighLevel}} \\
& \dot{Q}_{c,i} \geqslant 0
\end{aligned}
$$

Additionally, a rate of change penalty has been added to the inputs. The heat duty model constraint ($\dot{Q}_{c,i} = f(T_i, T_{\text{sp},i})$) is discussed in the next section.

A.3.2 Heat Duty Constraint

To obtain the heat duty model ($\dot{Q}_{c,i} = f(T_i, T_{\text{sp},i})$), we can break it up into two parts:

1. Model the regulatory temperature controller to determine the final controller

element signal as a function of the zone temperature and temperature setpoint

$$\nu_{air,i} = f(T_i, T_{sp,i})$$

2. Model the energy balance relating the final controller element to the actual heat duty

$$\dot{Q}_{c,i} = f(\nu_{air,i})$$

Both of these models can be identified from data. Since these models are often nonlinear in real buildings, assumptions will be made to form a linear model. Assuming a simple PI controller, the regulatory controller model is given by

$$\nu_{air,i} = \nu_{air,ss,i} + \tilde{K}_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right]$$
$$\varepsilon_i = T_{sp,i} - T_i$$

Next, the heat duty can be computed using the energy balance

$$\dot{Q}_{c,i} = U_i A_i (T_i - T_{CW})$$

where the heat transfer coefficient is nonlinear function

$$U = f(\nu_{air,i}, T_i, T_{CW}, HX_i)$$

Assuming a constant temperature gradient and a linear relationship between U and ν_{air} gives

$$\begin{aligned} U &\propto \nu_{air} \\ \dot{Q}_{c,i} = U_i A_i (T_i - T_{CW}) &\propto \nu_{air,i} \end{aligned}$$

Hence, the final simplified linear heat duty controller model is

$$\begin{aligned} \dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\int_0^t \varepsilon_i(t')dt'\right] \\ \varepsilon_i &= T_{sp,i} - T_i \end{aligned}$$

Saturation is represented as constraints on $\dot{Q}_{c,i}$ to keep the low-level airside problem as a linear program.

To enforce the high-level heat duty constraint along with cooling nonnegativity constraints, we replace the pointwise constraint

$$\begin{aligned} \sum_i \dot{Q}_{c,i} &\leqslant \dot{Q}_{HighLevel} \qquad \forall\, k \\ \dot{Q}_{c,i} &\geqslant 0 \qquad \forall\, k \end{aligned} \tag{A.1}$$

with

$$\begin{aligned} \frac{Q_{tot,k+1} - Q_{tot,k}}{\Delta} &\leqslant \dot{Q}_{HighLevel,k} \qquad \forall\, k \\ Q_{tot,k+1} - Q_{tot,k} &\geqslant 0 \qquad \forall\, k \end{aligned} \tag{A.2}$$

to obtain the final low-level airside MPC problem

$$
\begin{aligned}
\min_{\mathbf{x},\mathbf{u}} \quad & Q_{\text{total}} \\
\text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \quad \forall\, k \\
& \dot{Q}_{c,i} = \dot{Q}_{ss,i} + K_{c,i}\left[\varepsilon_i + \frac{1}{\tau_{I,i}}\bar{\varepsilon}_i\right] \\
& \frac{d\bar{\varepsilon}_i}{dt} = \varepsilon_i = T_{sp,i} - T_i \\
& \frac{dQ_{tot}}{dt} = \sum_i \dot{Q}_{c,i} \qquad \forall\, k \\
& T_{min} \leqslant T_i \leqslant T_{max} \qquad \forall\, k \\
& \frac{Q_{tot,k+1} - Q_{tot,k}}{\Delta} \leqslant \dot{Q}_{\text{HighLevel},k} \qquad \forall\, k \\
& Q_{tot,k+1} - Q_{tot,k} \geqslant 0 \qquad \forall\, k
\end{aligned}
$$

A.3.3 Formulation

The low-level problem is formulated as a linear program of the form:

$$
\min_z c^T z \qquad \text{s.t.} \quad Ez = e, \quad Fz \leqslant f, \quad l \leqslant z \leqslant u
$$

The subsequent sections outline the procedure for obtaining a linear program of this form.

Model

As shown in Table A.2, the states are the zone temperatures, the integrals of the zone tracking errors, and the total amount of cooling delivered. The inputs are the zone temperature setpoints. All states are assumed to be measured perfectly. The disturbances are the ambient temperature, external loads, and steady-state cooling rates.

States	$x = \begin{bmatrix} T_i & \bar{\varepsilon}_i & Q_{total} \end{bmatrix}^T$	$n = 2n_{bz} + 1$
Inputs	$u = \begin{bmatrix} T_{sp,i} \end{bmatrix}^T$	$m = n_{bz}$
Measurements	$y = \begin{bmatrix} T_i & Q_{total} \end{bmatrix}^T$	$p = n_{bz} + 1$
Disturbances	$d = \begin{bmatrix} T_a & \dot{Q}_{other,i} & \dot{Q}_{ss,i} \end{bmatrix}^T$	$n_d = 1 + 2n_{bz}$

Table A.2: Low-Level Problem Variables

The energy balance and heat duty model can be expressed as a continuous-time state-space model as shown on the following page. Discretize the continuous-time model above using the controller interval, Δ, to get the discrete version of the matrices: A, B, and B_d. Then, we can $C = I$ and $D = 0$ to get the final discrete-time state space model:

$$x^+ = Ax + Bu + B_d d$$

$$y = Cx + Du$$

$$\frac{d}{dt}\begin{bmatrix} T_1 \\ T_2 \\ \vdots \\ T_{n_{bz}} \\ \bar{\varepsilon}_1 \\ \bar{\varepsilon}_2 \\ \vdots \\ \bar{\varepsilon}_{n_{bz}} \\ Q_{tot} \end{bmatrix} = \quad \text{(continued on next page...)}$$

$$
\underbrace{\begin{bmatrix}
\frac{-H_1-\sum_{i\neq 1}\beta_{1i}+K_{c1}}{C_1} & \frac{\beta_{12}}{C_1} & \cdots & \frac{\beta_{1n_{bz}}}{C_1} & \frac{-K_{c1}}{C_1\tau_{l1}} & 0 & \cdots & 0 & 0 \\
\frac{\beta_{21}}{C_2} & \frac{-H_2-\sum_{i\neq 2}\beta_{2i}+K_{c2}}{C_2} & \cdots & \frac{\beta_{2n_{bz}}}{C_2} & 0 & \frac{-K_{c2}}{C_2\tau_{l2}} & \cdots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
\frac{\beta_{n_{bz}1}}{C_{n_{bz}}} & \frac{\beta_{n_{bz}2}}{C_{n_{bz}}} & \cdots & \frac{-H_{n_{bz}}-\sum_{i\neq n_{bz}}\beta_{n_{bz}i}+K_{cn_{bz}}}{C_{n_{bz}}} & 0 & 0 & \cdots & \frac{-K_{cn_{bz}}}{C_{n_{bz}}\tau_{ln_{bz}}} & 0 \\
-1 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 \\
0 & -1 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \cdots & -1 & 0 & 0 & \cdots & 0 & 0 \\
-K_{c,1} & -K_{c,2} & \cdots & -K_{c,n_{bz}} & \frac{K_{c1}}{\tau_{l1}} & \frac{K_{c2}}{\tau_{l2}} & \cdots & \frac{K_{cn_{bz}}}{\tau_{ln_{bz}}} & 0
\end{bmatrix}}_{A_c}
\begin{bmatrix} T_1 \\ T_2 \\ \vdots \\ T_{n_{bz}} \\ \bar{\varepsilon}_1 \\ \bar{\varepsilon}_2 \\ \vdots \\ \bar{\varepsilon}_{n_{bz}} \\ Q_{tot} \end{bmatrix}
$$

$$
+ \underbrace{\begin{bmatrix}
\frac{-K_{c,1}}{C_1} & 0 & \cdots & 0 \\
0 & \frac{-K_{c,2}}{C_2} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \frac{-K_{c,n_{bz}}}{C_{n_{bz}}} \\
1 & 0 & \cdots & 0 \\
0 & 1 & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & 1 \\
K_{c,1} & K_{c,2} & \cdots & K_{c,n_{bz}}
\end{bmatrix}}_{B_c}
\begin{bmatrix} T_{sp,1} \\ T_{sp,2} \\ \vdots \\ T_{sp,n_{bz}} \end{bmatrix}
$$

$$
+ \underbrace{\begin{bmatrix}
\frac{H_1}{C_1} & \frac{1}{C_1} & 0 & \cdots & 0 & \frac{-1}{C_1} & 0 & \cdots & 0 \\
\frac{H_2}{C_2} & 0 & \frac{1}{C_2} & \cdots & 0 & 0 & \frac{-1}{C_2} & \cdots & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\
\frac{H_{n_{bz}}}{C_{n_{bz}}} & 0 & 0 & \cdots & \frac{1}{C_{n_{bz}}} & 0 & 0 & \cdots & \frac{-1}{C_{n_{bz}}} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\
0 & 0 & 0 & \cdots & 0 & 1 & 1 & \cdots & 1
\end{bmatrix}}_{\Gamma_c}
\begin{bmatrix} T_a \\ \dot{Q}_{other,1} \\ \dot{Q}_{other,2} \\ \vdots \\ \dot{Q}_{other,n_{bz}} \\ \dot{Q}_{ss,1} \\ \dot{Q}_{ss,2} \\ \vdots \\ \dot{Q}_{ss,n_{bz}} \end{bmatrix}
$$

Variables

The variables in the linear program are

$$z = \begin{bmatrix} u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \end{bmatrix} \Bigg\} N(m+n) \quad \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_N \end{bmatrix} \Bigg\} Nn_{bz} \quad \begin{bmatrix} p_{Q,1} \\ p_{Q,2} \\ \vdots \\ p_{Q,N} \end{bmatrix} \Bigg\} N \quad \begin{bmatrix} p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \end{bmatrix} \Bigg\} Nm$$

The total number of variables is $N(2m + n + n_{bz} + 1)$ with rate of change penalties and the total number of variables is $N(m + n + n_{bz} + 1)$ without rate of change penalties.

Objective

The linear objective function is

$$
c = \begin{matrix} u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \\ p_1 \\ p_2 \\ \vdots \\ p_N \\ p_{Q,1} \\ p_{Q,2} \\ \vdots \\ p_{Q,N} \\ p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \end{matrix}
\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ c_Q \\ \beta \\ \beta \\ \vdots \\ \beta \\ \gamma \\ \gamma \\ \vdots \\ \gamma \\ \alpha \\ \alpha \\ \alpha \\ \vdots \\ \alpha \end{bmatrix}
\qquad
c_Q = \begin{matrix} x_1(N) \\ x_2(N) \\ \vdots \\ x_{n-1}(N) \\ x_n(N) \end{matrix}
\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}
$$

Constraints

The state equation ($x^+ = Ax + Bu + B_d d$) is represented as $Ez = e$ where

$$E = \begin{array}{c} \begin{array}{cccccccccccc} u(0) & x(1) & u(1) & x(2) & u(2) & x(3) & \cdots & u(N-1) & x(N) & p_1 & \cdots \end{array} \\ \begin{bmatrix} -B & I & & & & & & & & & \\ & -A & -B & I & & & & & & & \\ & & & -A & -B & I & & & & & \\ & & & & & & \ddots & & & & \\ & & & & & & -A & -B & I & & \end{bmatrix} \end{array}$$

$$e = \begin{bmatrix} B_d d(0) \\ B_d d(1) \\ B_d d(2) \\ \vdots \\ B_d d(N-1) \end{bmatrix} + \left.\begin{bmatrix} Ax(0) \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}\right\} Nn$$

The high-level heat duty constraint ($\frac{Q_{tot,k+1} - Q_{tot,k}}{\Delta} \leqslant \dot{Q}_{HighLevel,k}$) is represented as $F_1 z \leqslant f_1$ where

$$F_1 = \begin{array}{c} \begin{array}{cccccccccccccc} u(0) & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & x(N) & \cdots & p_{Q,1} & p_{Q,2} & p_{Q,3} & \cdots & p_{Q,N} & \cdots \end{array} \\ \begin{bmatrix} & V & & & & & & & & -I & & & & & \\ & -V & & V & & & & & & & -I & & & & \\ & & & -V & & V & & & & & & -I & & & \\ & & & & & \ddots & & & & & & & \ddots & & \\ & & & & & -V & & V & & & & & & -I & \end{bmatrix} \end{array}$$

$$f_1 = \begin{bmatrix} \dot{Q}_{\text{HighLevel}}(0) \\ \dot{Q}_{\text{HighLevel}}(1) \\ \dot{Q}_{\text{HighLevel}}(2) \\ \vdots \\ \dot{Q}_{\text{HighLevel}}(N-1) \end{bmatrix} \Delta + \left.\begin{bmatrix} Q_{\text{tot}}(0) \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}\right\} N$$

and $V = [0 \quad 0 \quad \cdots \quad 0 \quad 1]$ is a row vector of length n.

The nonnegativity constraint ($Q_{\text{tot},k+1} - Q_{\text{tot},k} \geqslant 0$) is represented as $F_2 z \leqslant f_2$ where

$$F_2 = \begin{array}{c} \begin{matrix} u(0) & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & x(N) & \cdots \end{matrix} \\ \begin{bmatrix} & -V & & & & & & & \\ & V & & -V & & & & & \\ & & & V & & -V & & & \\ & & & & & \ddots & & & \\ & & & & & V & & -V & \end{bmatrix} \end{array}$$

$$f_2 = \left.\begin{bmatrix} -Q_{\text{tot}}(0) \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}\right\} N$$

The rate of change constraint, if included, is represented as $F_3 z \leqslant f_3$ where

$$
F_3 = \begin{array}{c}
\begin{array}{ccccccccccccccc}
u(0) & & x(1) & u(1) & x(2) & u(2) & \cdots & u(N-1) & \cdots & p_{0,-1} & p_{1,0} & p_{2,1} & \cdots & p_{N-1,N-2} &
\end{array} \\
\begin{bmatrix}
I & & & & & & & & & -I & & & & & \\
-I & & & & & & & & & -I & & & & & \\
-I & & & I & & & & & & & -I & & & & \\
I & & & -I & & & & & & & -I & & & & \\
 & & & -I & & I & & & & & & -I & & & \\
 & & & I & & -I & & & & & & -I & & & \\
 & & & & & & \ddots & & & & & & \ddots & & \\
 & & & & & -I & & I & & & & & & -I & \\
 & & & & & I & & -I & & & & & & -I &
\end{bmatrix}
\end{array}
$$

$$
f_3 = \left.\begin{bmatrix}
u(-1) \\
-u(-1) \\
0 \\
0 \\
0 \\
0 \\
\vdots \\
0 \\
0
\end{bmatrix}\right\} 2Nm
$$

The comfort zone constraint ($T_{min} \leqslant T_i \leqslant T_{max}$) is represented as $F_4 z \leqslant f_4$ where

$$
F_4 = \begin{array}{c}
\begin{array}{ccccccccccc}
u(0) & x(1) & u(1) & x(2) & \cdots & u(N-1) & x(N) & p_1 & p_2 & \cdots & p_N \quad \cdots
\end{array} \\
\begin{bmatrix}
 & -G & & & & & & -I & & & & \\
 & G & & & & & & -I & & & & \\
 & & & -G & & & & & -I & & & \\
 & & & G & & & & & -I & & & \\
 & & & & \ddots & & & & & \ddots & & \\
 & & & & & & -G & & & & -I & \\
 & & & & & & G & & & & -I &
\end{bmatrix}
\end{array}
$$

$$
f_4 = \left.\begin{bmatrix} -T_{min} \\ T_{max} \\ -T_{min} \\ T_{max} \\ \vdots \\ -T_{min} \\ T_{max} \end{bmatrix}\right\} 2Nn_{bz}
$$

and where $G = [I \quad 0]$, T_{min} is a column vector of length z with the lower end of the temperature comfort zone, and T_{max} is a column vector of length z with the upper end of the temperature comfort zone.

The bounds on the variables are

$$
l = \begin{bmatrix} u_{min} \\ x_{min} \\ u_{min} \\ x_{min} \\ \vdots \\ u_{min} \\ x_{min} \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \leqslant \begin{bmatrix} u(0) \\ x(1) \\ u(1) \\ x(2) \\ \vdots \\ u(N-1) \\ x(N) \\ p_1 \\ p_2 \\ \vdots \\ p_N \\ p_{Q,1} \\ p_{Q,2} \\ \vdots \\ p_{Q,N} \\ p_{0,-1} \\ p_{1,0} \\ p_{2,1} \\ \vdots \\ p_{N-1,N-2} \end{bmatrix} \leqslant \begin{bmatrix} u_{max} \\ x_{max} \\ u_{max} \\ x_{max} \\ \vdots \\ u_{max} \\ x_{max} \\ \infty \\ \infty \\ \vdots \\ \infty \\ \infty \\ \infty \\ \vdots \\ \infty \\ \infty \\ \infty \\ \infty \\ \vdots \\ \infty \end{bmatrix} = u
$$

where

$$u_{min} = \begin{bmatrix} -\infty \\ \vdots \\ -\infty \end{bmatrix} \leqslant \begin{bmatrix} T_{sp,1} \\ \vdots \\ T_{sp,n_{bz}} \end{bmatrix} \leqslant \begin{bmatrix} \infty \\ \vdots \\ \infty \end{bmatrix} = u_{max}$$

$$x_{min} = \begin{bmatrix} -\infty \\ \vdots \\ -\infty \\ -\infty \\ \vdots \\ -\infty \\ -\infty \end{bmatrix} \leqslant \begin{bmatrix} T_1 \\ \vdots \\ T_{n_{bz}} \\ \bar{\varepsilon}_1 \\ \vdots \\ \bar{\varepsilon}_{n_{bz}} \\ Q_{tot} \end{bmatrix} \leqslant \begin{bmatrix} \infty \\ \vdots \\ \infty \\ \infty \\ \vdots \\ \infty \\ \infty \end{bmatrix} = x_{max}$$

By defining

$$F = \begin{bmatrix} F_1 \\ F_2 \\ F_3 \\ F_4 \end{bmatrix} \qquad f = \begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{bmatrix}$$

these constraints can be represented in the form

$$Ez = e$$

$$Fz \leqslant f$$

$$l \leqslant z \leqslant u$$

BIBLIOGRAPHY

[1] Abraxas Energy. Typical HVAC Rules of Thumb. http://www.abraxasenergy.com/energy-resources/toolbox/hvac-rules-of-thumb/, May 2016.

[2] A. Afram and F. Janabi-Sharifi. Theory and applications of HVAC control systems—A review of model predictive control (MPC). *Build. Environ.*, 72: 343–355, Feb 2014.

[3] R. Amrit, J. B. Rawlings, and D. Angeli. Economic optimization using model predictive control with a terminal cost. *Annual Rev. Control*, 35:178–186, 2011.

[4] M. Avci, M. Erkoc, A. Rahmani, and S. Asfour. Model predictive HVAC load control in buildings using real-time electricity pricing. *Energ. Buildings*, 60: 199–209, May 2013.

[5] T. N. Aynur. Variable refrigerant flow systems: A review. *Energ. Buildings*, 42 (7):1106–1112, 2010.

[6] M. Aziz, T. Oda, T. Mitani, Y. Watanabe, and T. Kashiwagi. Utilization of electric vehicles and their used batteries for peak-load shifting. *Energies*, 8(5): 3720–3738, 2015.

[7] S. C. Bengea, A. D. Kelman, F. Borrelli, R. Taylor, and S. Narayanan. Implementation of model predictive control for an HVAC system in a mid-size commercial building. *HVAC&R Res.*, 20:121–135, 2014.

[8] G. Berckmans, M. Messagie, J. Smekens, N. Omar, L. Vanhaverbeke, and J. Van Mierlo. Cost projection of state of the art lithium-ion batteries for electric vehicles up to 2030. *Energies*, 10(9), 2017.

[9] S. Blair. Editors' choice and best energy/industrial: Stanford Energy System Innovations. *Engineering News-Record*, 2016. URL http://www.enr.com/articles/39005-editors-choice-best-energyindustrial-stanford-energy-system-innovations.

[10] J. E. Braun and G. T. Diderrich. Near-optimal control of cooling towers for chilled-water systems. *ASHRAE Transactions*, 96(CONF-9006117), 1990.

[11] J. Cai, D. Kim, R. Jaramillo, J. Braun, and J. Hu. A general multi-agent control approach for building energy system optimization. *Energ. Buildings*, 127:337–351, 2016.

[12] P. Christofides, R. Scattolini, D. de la Peña, and J. Liu. Distributed model predictive control: A tutorial review and future research directions. *Comput. Chem. Eng.*, 51:21–41, 2013.

[13] W. J. Cole, T. F. Edgar, and A. Novoselac. Use of model predictive control to enhance the flexibility of thermal energy storage cooling systems. In *American Control Conference (ACC), 2012*, pages 2788–2793, 2012.

[14] A. J. del Real, A. Arce, and C. Bordons. Combined environmental and economic dispatch of smart grids using distributed model predictive control. *Int. J. Elec. Pow. & Energ. Sys.*, 54:65–76, Jan 2014.

[15] Department of Energy. *2011 Buildings Energy Data Book*. Mar 2012. `http://buildingsdatabook.eren.doe.gov/ChapterIntro3.aspx`.

[16] M. Diehl, R. Amrit, and J. B. Rawlings. A Lyapunov function for economic optimizing model predictive control. *IEEE Trans. Auto. Cont.*, 56(3):703–707, 2011.

[17] K. Divya and J. Østergaard. Battery energy storage technology for power systems–an overview. *Electr. Pow. Sys. Res.*, 79(4):511–520, 2009.

[18] A. Dounis and C. Caraiscos. Advanced control systems engineering for energy and comfort management in a building environment—A review. *Ren. Sust. Energ. Rev.*, 13:1246–1261, 2009.

[19] A. W. Dowling, R. Kumar, and V. M. Zavala. A multi-scale optimization framework for electricity market participation. *Appl. Energ.*, 190:147–164, 2017.

[20] M. Elliott and B. Rasmussen. Neighbor-communication model predictive control and HVAC systems. In *American Control Conference*, pages 3020–3025, Montreal, Canada, June 27–29, 2012.

[21] M. Ellis, H. Durand, and P. D. Christofides. A tutorial review of economic model predictive control methods. *J. Proc. Cont.*, 24(8):1156–1178, 2014.

[22] M. Ellis, J. Liu, and P. D. Christofides. *Economic Model Predictive Control: Theory, Formulations and Chemical Process Applications*. Springer, London, England, 2016.

[23] Energy Information Administration. Energy Consumption Surveys. `https://www.eia.gov/tools/faqs/faq.php?id=86`, May 2018.

[24] J. M. Fenton. Electric vehicles will save the world. *Electrochem. Soc. Interface*, 25(2):29–32, 2016.

[25] A. Ferramosca, D. Limon, A. H. Gonzalez, D. Odloak, and E. F. Camacho. MPC for tracking zone regions. *Journal of Process Control*, 20(4):506–516, 2010.

[26] R. Z. Freire, G. H. Oliveira, and N. Mendes. Predictive controllers for thermal comfort optimization and energy savings. *Energ. Buildings*, 40:1353–1365, 2008.

[27] R. Goebel, R. G. Sanfelice, and A. R. Teel. Hybrid dynamical systems. *IEEE Ctl. Sys. Mag.*, pages 28–93, April 2009.

[28] W. Guo and M. Zhou. Technologies toward thermal comfort-based and energy-efficient HVAC systems: A review. In *Systems, Man and Cybernetics, 2009. SMC 2009. IEEE International Conference on*, pages 3883–3888, October 2009.

[29] A. Hajiah and M. Krarti. Optimal controls of building storage systems using both ice storage and thermal mass—Part II: Parametric analysis. *Energ. Convers. Manage.*, 64(0):509–515, 2012.

[30] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen. Economic model predictive control for building climate control in a smart grid. In *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, pages 1–6, Washington, DC, USA, 2012.

[31] M. F. Haniff, H. Selamat, R. Yusof, S. Buyamin, and F. S. Ismail. Review of HVAC scheduling techniques for buildings towards energy-efficient and cost-effective operations. *Ren. Sust. Energ. Rev.*, 27:94–103, 2013.

[32] G. P. Henze. Energy and cost minimal control of active and passive building thermal storage inventory. *J. Solar Ener. Eng.*, 127(3):343–351, 2005.

[33] G. P. Henze, C. Felsmann, and G. Knabe. Evaluation of optimal control for active and passive building thermal storage. *Int. J. Thermal Sci.*, 43(2):173–183, 2004.

[34] C. Heymans, S. B. Walker, S. B. Young, and M. Fowler. Economic analysis of second use electric vehicle batteries for residential energy storage and load-levelling. *Energy Policy*, 71:22–30, 2014.

[35] W. Huang and H. Lam. Using genetic algorithms to optimize controller parameters for HVAC systems. *Energ. Buildings*, 26:277–282, 1997.

[36] G.-Y. Jin, W.-J. Cai, L. Lu, E. L. Lee, and A. Chiang. A simplified modeling of mechanical cooling tower for control and optimization of HVAC systems. *Energ. Convers. Manage.*, 48(2):355–365, 2007.

[37] S. A. Kalogirou. Artificial neural networks in renewable energy systems applications: A review. *Ren. Sust. Energ. Rev.*, 5:373–401, 2001.

[38] S. A. Kalogirou. Artificial neural networks and genetic algorithms in energy applications in buildings. *Adv. Build. Energ. Res.*, 3:83–119, 2009.

[39] C. M. Kellett. A compendium of comparison function results. *Math. Contr. Sign. Syst.*, 26(3):339–374, 2014. doi: 10.1007/s00498-014-0128-8.

[40] A. Kelman, Y. Ma, and F. Borrelli. Analysis of local optima in predictive control for energy efficient buildings. *J. Build. Perf. Sim.*, pages 1–20, May 2012.

[41] M. Kintner-Meyer and A. F. Emery. Optimal control of an HVAC system using cold storage and building thermal capacitance. *Energ. Buildings*, 23(1):19–31, 1995.

[42] S. Koehler and F. Borrelli. Building temperature distributed control via explicit MPC and "trim and respond" methods. In *European Control Conference*, pages 4334–4339, Zurich, Switzerland, 2013.

[43] D. Kolokotsa, D. Tsiavos, G. Stavrakakis, K. Kalaitzakis, and E. Antonidakis. Advanced fuzzy logic controllers design and evaluation for buildings' occupants thermal-visual comfort and indoor air quality satisfaction. *Energ. Buildings*, 33:531–543, 2001.

[44] R. Kumar, M. J. Wenzel, M. J. Ellis, M. N. ElBsat, K. H. Drees, and V. M. Zavala. A stochastic model predictive control framework for stationary battery systems. *IEEE Trans. Power Systems*, 33(4):4397–4406, Jul 2018.

[45] A. Kusiak and G. Xu. Modeling and optimization of HVAC systems using a dynamic neural network. *Energy*, 42:241–250, 2012. 8th World Energy System Conference, WESC 2010.

[46] M. Y. Lamoudi, M. Alamir, and P. Béguery. Distributed constrained model predictive control based on bundle method for building energy management. In *50th IEEE Conference on Decision and Control and European Control Conference*, pages 8118–8124, Orlando, FL, USA, 2011.

[47] M. Y. Lamoudi, M. Alamir, and P. Béguery. Model predictive control for energy management in buildings. Part 2 : Distributed model predictive control. In *IFAC workshop on Nonlinear Model Predictive Control*, pages 226–231, 2012.

[48] M. Y. Lamoudi, M. Alamir, and P. Béguery. Model predictive control for energy management in buildings. Part 1 : Zone model predictive control. In *IFAC workshop on Nonlinear Model Predictive Control*, pages 21–26, 2012.

[49] O. Lavrova, F. Cheng, S. Abdollahy, H. Barsun, A. Mammoli, D. Dreisigmayer, S. Willard, B. Arellano, and C. van Zeyl. Analysis of battery storage utilization for load shifting and peak smoothing on a distribution feeder in new mexico. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, pages 1–6, Jan 2012.

[50] T.-S. Lee, K.-Y. Liao, and W.-C. Lu. Evaluation of the suitability of empirically-based models for predicting energy performance of centrifugal water chillers with variable chilled water flow. *Appl. Energ.*, 93(0):583–595, 2012.

[51] L. Ljung. *System Identification - Theory for the User*. Prentice Hall, New Jersey, second edition, 1999.

[52] J. Ma, J. Qin, T. Salsbury, and P. Xu. Demand reduction in building energy systems based on economic model predictive control. *Chem. Eng. Sci.*, 67(1): 92 – 100, 2012.

[53] Y. Ma, G. Anderson, and F. Borrelli. A distributed predictive control approach to building temperature regulation. In *American Control Conference (ACC), 2011*, pages 2089–2094. IEEE, 2011.

[54] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves. Model predictive control for the operation of building cooling systems. *IEEE Ctl. Sys. Tech.*, 20(3):796–803, 2012.

[55] Y. Ma, A. Kelman, A. Daly, and F. Borrelli. Predictive control for energy efficient buildings with thermal storage. *IEEE Ctl. Sys. Mag.*, pages 44–64, Feb 2012.

[56] Y. Ma, S. Richter, and F. Borrelli. Chapter 14: Distributed model predictive control for building temperature regulation. In *Control and Optimization with Differential-Algebraic Constraints*, volume 22, chapter 14, pages 293–314. SIAM, March 2012.

[57] B. Mayer, M. Killian, and M. Kozek. Management of hybrid energy supply systems in buildings using mixed-integer model predictive control. *Energ. Convers. Manage.*, 98:470–483, 2015.

[58] D. I. Mendoza-Serrano and D. J. Chmielewski. HVAC control using infinite-horizon economic MPC. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 6963–6968, 2012.

[59] P.-D. Moroşan, R. Bourdais, D. Dumur, and J. Buisson. Building temperature regulation using a distributed model predictive control. *Energ. Buildings*, 42: 1445–1452, Sep 2010.

[60] P.-D. Moroşan, R. Bourdais, D. Dumur, and J. Buisson. A distributed MPC strategy based on Benders' decomposition applied to multi-source multi-zone temperature regulation. *J. Proc. Cont.*, 21:729–737, Jun 2011.

[61] M. Mossolly, K. Ghali, and N. Ghaddar. Optimal control strategy for a multi-zone air conditioning system using a genetic algorithm. *Energy*, 34:58–66, 2009.

[62] B. Nykvist and M. Nilsson. Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5:329–332, Mar 2015.

[63] B. J. Odelson, M. R. Rajamani, and J. B. Rawlings. A new autocovariance least-squares method for estimating noise covariances. *Automatica*, 42(2):303–308, February 2006.

[64] F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *American Control Conference (ACC), 2010*, pages 5100–5105, Baltimore, MD, USA, 2010.

[65] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. 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*, pages 1927–1932, Atlanta, GA, USA, December 2010.

[66] N. R. Patel and J. B. Rawlings. Model predictive control for optimal zone tracking. AIChE Annual Meeting, San Francisco, CA, Nov 2016.

[67] N. R. Patel and J. B. Rawlings. Economic optimization of large-scale embedded battery applications. AIChE Annual Meeting, Minneapolis, MN, Nov 2017.

[68] N. R. Patel and J. B. Rawlings. Model predictive control for tracking zones with discrete actuators. AIChE Midwest Regional Conference, Chicago, IL, Mar 2017.

[69] N. R. Patel and J. B. Rawlings. Applications of MPC to building HVAC systems. In S. V. Raković and W. S. Levine, editors, *Handbook of Model Predictive Control*, chapter 3.5. Springer International Publishing, 2018.

[70] N. R. Patel, M. J. Risbeck, J. B. Rawlings, M. J. Wenzel, and R. D. Turney. Distributed economic model predictive control for large-scale building temperature regulation. In *American Control Conference*, pages 895–900, Boston, MA, July 6–8, 2016.

[71] N. R. Patel, J. B. Rawlings, M. J. Ellis, M. J. Wenzel, and R. D. Turney. An economic model predictive control framework for distributed embedded battery applications. In *5th International High Performance Buildings Conference at Purdue*, West Lafayette, IN, July 9–12, 2018.

[72] N. R. Patel, M. J. Risbeck, J. B. Rawlings, C. T. Maravelias, M. J. Wenzel, and R. D. Turney. A case study of economic optimization of HVAC systems based on the Stanford University campus airside and waterside systems. In *5th International High Performance Buildings Conference at Purdue*, West Lafayette, IN, July 9–12, 2018.

[73] N. R. Patel, J. B. Rawlings, M. J. Ellis, M. J. Wenzel, and R. D. Turney. Economic optimization of distributed embedded battery units for large-scale HVAC applications. *AIChE J.*, submitted 2018. Special PSE Issue on Sustainable Energy.

[74] N. R. Patel, M. J. Risbeck, J. B. Rawlings, C. T. Maravelias, M. J. Wenzel, and R. D. Turney. A case study of central plant and airside optimization based on the stanford university campus. *Energ. Buildings*, in prep 2018.

[75] D. Picard, M. Sourbron, F. Jorissen, Z. Váňa, J. Cigler, L. Ferkl, and L. Helsen. Comparison of model predictive control performance using grey-box and white-box controller models of a multi-zone office building. In *4th International High Performance Buildings Conference at Purdue*, West Lafayette, IN, July 11–14, 2016.

[76] D. Picard, J. Drgoňa, M. Kvasnica, and L. Helsen. Impact of the controller model complexity on model predictive control performance for buildings. *Energ. Buildings*, 152:739–751, Oct 2017.

[77] V. Putta, G. Zhu, D. Kim, J. Hu, and J. E. Braun. A distributed approach to efficient model predictive control of building HVAC systems. In *International High Performance Buildings Conference*, pages 1–10, 2012.

[78] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Eng. Pract.*, 11(7):733–764, 2003.

[79] J. B. Rawlings and R. Amrit. Optimizing process economic performance using model predictive control. In L. Magni, D. M. Raimondo, and F. Allgöwer, editors, *Nonlinear Model Predictive Control*, volume 384 of *Lecture Notes in Control and Information Sciences*, pages 119–138, Berlin, 2009. Springer.

[80] J. B. Rawlings and D. Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, Madison, WI, 2009. 668 pages, ISBN 978-0-9759377-0-9.

[81] J. B. Rawlings and M. J. Risbeck. Model predictive control with discrete actuators: Theory and application. *Automatica*, 78:258–265, 2017.

[82] J. B. Rawlings and B. T. Stewart. Coordinating multiple optimization-based controllers: New opportunities and challenges. *J. Proc. Cont.*, 18:839–845, 2008.

[83] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Design, and Computation*. Nob Hill Publishing, Madison, WI, 2nd edition, 2017. 770 pages, ISBN 978-0-9759377-3-0.

[84] J. B. Rawlings, N. R. Patel, M. J. Risbeck, C. T. Maravelias, M. J. Wenzel, and R. D. Turney. Economic MPC and real-time decision making with application to large-scale HVAC energy systems. *Comput. Chem. Eng.*, 114:89–98, 2018. FOCAPO/CPC 2017.

[85] M. J. Risbeck, C. T. Maravelias, J. B. Rawlings, and R. D. Turney. Cost optimization of combined building heating/cooling equipment via mixed-integer linear programming. In *American Control Conference*, pages 1689–1694, Chicago, IL, July 1–3, 2015.

[86] M. J. Risbeck, C. T. Maravelias, J. B. Rawlings, and R. D. Turney. A mixed-integer linear programming model for real-time cost optimization of building heating, ventilation, and air conditioning equipment. *Energ. Buildings*, 142: 220–235, 2017.

[87] T. I. Salsbury. A survey of control technologies in the building automation industry. In *16th IFAC World Congress*, pages 1396–1407, 2005.

[88] H. Scherer, M. Pasamontes, J. Guzmán, J. Álvarez, E. Camponogara, and J. Normey-Rico. Efficient building energy management using distributed model predictive control. *J. Proc. Cont.*, 24(6):740–749, June 2014.

[89] P. O. M. Scokaert and J. B. Rawlings. Feasibility issues in linear model predictive control. *AIChE J.*, 45(8):1649–1659, August 1999.

[90] J. Singh, N. Singh, and J. Sharma. Fuzzy modeling and control of HVAC systems—A review. *J. Sci. Ind. Res.*, 65:470, 2006.

[91] S. Soyguder, M. Karakose, and H. Alli. Design and simulation of self-tuning PID-type fuzzy adaptive control for an expert HVAC system. *Exp. Sys. App.*, 36:4566–4573, 2009.

[92] J. Stagner. Enterprise optimization solution (EOS) cost savings vs. manual plant dispatching. Report on Central Energy Facility, Stanford Energy System Innovations, 2016.

[93] B. T. Stewart, A. N. Venkat, J. B. Rawlings, S. J. Wright, and G. Pannocchia. Cooperative distributed model predictive control. *Sys. Cont. Let.*, 59:460–469, 2010.

[94] K. Subramanian, C. T. Maravelias, and J. B. Rawlings. A state-space model for chemical production scheduling. *Comput. Chem. Eng.*, 47:97–110, December 2012.

[95] C. R. Touretzky and M. Baldea. Nonlinear model reduction and model predictive control of residential buildings with energy recovery. *J. Proc. Cont.*, pages 1–17, Nov 2013.

[96] C. R. Touretzky and M. Baldea. A hierarchical scheduling and control strategy for thermal energy storage systems. *Energ. Buildings*, 110:94–107, 2016.

[97] S. Wang and Z. Ma. Supervisory and optimal control of building HVAC systems: A review. *HVAC&R Res.*, 14:3–32, 2008.

[98] M. A. Zagrobelny and J. B. Rawlings. Practical improvements to autocovariance least-squares. *AIChE J.*, 61:1840–1855, 2015.

[99] B. Zakeri and S. Syri. Electrical energy storage systems: A comparative life cycle cost analysis. *Renewable and Sustainable Energy Reviews*, 42:569–596, 2015.

[100] V. M. Zavala. A multiobjective optimization perspective on the stability of economic MPC. *IFAC–P. Online*, 48(8):974–980, 2015.

[101] V. M. Zavala, E. M. Constantinescu, T. Krause, and M. Anitescu. On-line economic optimization of energy systems using weather forecast information. *J. Proc. Cont.*, 19:1725–1736, Dec 2009.

INDEX