DECENTRALIZED AND CENTRALIZED OPTIMAL CONTROL
OF HVAC SYSTEMS

by
Ryan S. Montrose

A dissertation
submitted in partial fulfillment
of the requirements for the degree of
Master of Science in Mechanical Engineering
Boise State University
August 2021

© 2021

Ryan S. Montrose

ALL RIGHTS RESERVED

BOISE STATE UNIVERSITY GRADUATE COLLEGE

DEFENSE COMMITTEE AND FINAL READING APPROVALS

of the **dissertation** submitted by

Ryan S. Montrose

Dissertation Title: Decentralized and Centralized Optimal Control of HVAC Systems

Date of Final Oral Examination: 30 April 2021

The following individuals read and discussed the dissertation submitted by student Ryan S. Montrose, and they evaluated the student's presentation and response to questions during the final oral examination. They found that the student passed the final oral examination.

John Gardner Ph.D., P.E.	Chair, Supervisory Committee
--------------------------	------------------------------

Aykut C. Satici Ph.D.	Co-Chair, Supervisory Committee
-----------------------	---------------------------------

John Chaisson Ph.D.	Member, Supervisory Committee
---------------------	-------------------------------

Zhangxian Deng Ph.D.	Member, Supervisory Committee
----------------------	-------------------------------

The final reading approval of the **dissertation** was granted by John Gardner Ph.D., P.E., Chair of the Supervisory Committee. The **dissertation** was approved by the Graduate College.

DEDICATION

To Joyce and Lewis Montrose, the ones who opened the door to higher education,
and a brighter future.

ACKNOWLEDGMENT

I would like to personally thank John Gardner for his keen insight in both myself and my research, and Aykut Satici for his tireless help. In addition, I would like to thank my committee members, John Chaisson and Zhangxian Deng, whom have developed my skills as an engineer and provided their expertise with my research. Lastly I would like to thank my peers who have helped me along this journey.

ABSTRACT

Utility service providers are often challenged with the synchronization of thermostatically controlled loads. Load synchronization, resulting from naturally occurring or demand response events, can damage power distribution equipment and reduce the grid's efficiency. Because thermostatically controlled loads constitute most of the power consumed by the grid at any given time, the proper control of such devices can lead to significant energy savings and improved grid stability. The contribution of this thesis is developing optimal control algorithms for both single-stage and variable-speed heat pump HVAC systems. Our control architecture allows for regulating home temperatures through selective peer-to-peer communication while simultaneously minimizing aggregate power consumption and aggregate load volatility. For comparison purposes, various low-level and centralized optimal controllers are explored and compared against their decentralized counterparts.

TABLE OF CONTENTS

DEDICATION	iv
ACKNOWLEDGMENT	v
ABSTRACT	vi
LIST OF FIGURES	x
LIST OF TABLES	xii
LIST OF ABBREVIATIONS	xiii
1 INTRODUCTION	1
1.1 History	1
1.2 Motivation	2
2 BACKGROUND	10
2.1 Modeling Approach	11
2.1.1 SSHP Dynamic Model	15
2.1.2 VSHP Dynamic Model	17
2.2 Network Communication	18
2.3 Baseline Control Frameworks	21

2.4	Parameter Estimation	25
2.5	SSHP Lockout Procedure	30
2.6	Demand Response	32
3	CONTROLLER	36
3.1	Decentralized Discrete Controller (DD)	39
3.2	Centralized Discrete Controller (CD)	44
3.3	Decentralized Continuous Controller (DC)	47
3.4	Centralized Continuous Controller (CC)	50
4	CASE STUDY	53
4.1	Baseline Simulation Results	55
4.1.1	Hysteresis Controller	56
4.1.2	Advanced Hysteresis Controller	58
4.1.3	PID Controller	60
4.2	Optimal Control Simulation Results	62
4.2.1	Decentralized Discrete Controller	63
4.2.2	Centralized Discrete Controller	66
4.2.3	Decentralized Continuous Controller	69
4.2.4	Centralized Continuous Controller	72
4.3	Comparison	75
5	CONCLUSION	81
	REFERENCES	82
	APPENDICES	86

A	DERIVATIONS	87
A.1	Second-order Equivalent Thermal Parameter (ETP) model	88

LIST OF FIGURES

2.1	Second-order equivalent thermal parameter circuit diagram.	12
2.2	Single Speed Heat Pump (SSHP), $m \in \{0, 1\}$, and Variable Speed Heat Pump (VSHP), $m \in [0, 1]$, control states, where the SSHP is either Off or ON and VSHP is linearly between OFF and ON.	15
2.3	Unweighted random regular graph ($N = 50$, $N_{cd} = 4$)	20
2.4	Closed loop controller for VSHP.	23
3.1	Four control methods explored, where $\mathbf{M} = [m_1, \dots, m_N]^\top$	37
3.2	System Diagram.	38
3.3	Objective soft constraint $\tilde{T}$ calculated using double hinge function. . .	43
4.1	Hysteresis controller results generated via Algorithm (2) (N=1,000). .	57
4.2	Advanced hysteresis controller results generated via Algorithm (3) (N=1,000). .	59
4.3	Advanced hysteresis controller results generated via Algorithm (4) (N=1,000). .	61
4.4	Decentralized Discrete (DD) controller results generated via Algorithm (5) (N=50).	64
4.5	DD estimated thermal parameter values represented as error signal (N=50).	65
4.6	Centralized Discrete (CD) controller results generated via Algorithm (6) (N=50).	67

4.7	CD estimated thermal parameter values represented as error signal (N=50).	68
4.8	Decentralized Continuous (DC) controller results generated via Algo- rithm (7) (N=50).	70
4.9	DC estimated thermal parameter values represented as error signal (N=50).	71
4.10	Centralized Continuous (CC) controller results generated via Algo- rithm (8) (N=50).	73
4.11	CC estimated thermal parameter values represented as error signal (N=50).	74

LIST OF TABLES

4.1	Thermal Parameters.	55
4.2	Simulation parameters	55
4.3	Quantitative Performance Metrics.	76

LIST OF ABBREVIATIONS

CC Centralized Continuous

CD Centralized Discrete

DC Decentralized Continuous

DD Decentralized Discrete

ETP Equivalent Thermal Parameter

FFT Fast Fourier Transform

HVAC Heating Ventilation and Air Conditioning

MIQP Mixed Integer Quadratic Program

MPC Model Predictive Control

NSRDB National Solar Radiation Database

P2P Peer-to-Peer

PID Proportional Integral Derivative

QP Quadratic Program

RLS Recursive Least Squares

SSHP Single Speed Heat Pump

TCL Thermostatically Controlled Load

TMY Typical meteorological Year

VSHP Variable Speed Heat Pump

CHAPTER 1:

INTRODUCTION

1.1 History

From its inception to our modern-day consumption needs, our electric utilities have experienced a series of technological advancements, thereby making power more accessible to the public. Beginning in the late 19th century, inventor and businessman, Thomas Edison developed the world's first commercially available incandescent lightbulb. Out of the necessity to power Edison's invention, various electric utilities were created.

Power generating stations like Pearl Street Station in New York City quickly popped up in dense urban sprawls. The general belief that electric utilities were unsafe and inefficient started to change after several advancements such as the alternating power sources, introduced by Nikola Tesla. Between Edison and Tesla's era to before the onset of World War I, electric utilities experienced steady growth. Within this period, developments in the generation and distribution of power brought electricity to many homes where entry barriers became less prominent for the public. After World War II, an explosion in consumer and industrial-grade electronics flooded the market. Everything from Heating Ventilation and Air Conditioning (HVAC) systems to cooking appliances, and all things in between, began to saturate the demand for

electricity. Because of this large influx of demanded power, electric utilities also expanded to capture this emergent market. New energy sources were brought to market, like nuclear fission and hydroelectricity, thereby softening the burden placed on fossil fuels. Historically, this moment is seen by many as the birth of America's first smart grid system (Johnson, 2010).

Much like the necessity for horizontal integration to meet up with demand, these utilities also needed to balance various energy sources. In the Cold War era, coincident to the computer boom of the late 20th century, the aid of sophisticated control systems began to revolutionize smart grid systems. During this computer boom, significant advancements, brought forth by modern control theory, greatly improved power distribution's scope and efficiency. From the start of the Cold War to the present day, many industrial countries began reducing their fossil fuel usage. Many renewable energy sources have penetrated the market space. Such renewable sources include wind, solar, geothermal, and the continual development of hydroelectric reservoirs. Nuclear fission also expanded during this period, but public support often varies in response to natural disasters. As of recently, many public and private ventures are researching the feasibility of fusion reactors.

1.2 Motivation

Although high penetration levels in renewable energy sources are observed as a practical step in curtailing fossil fuel usage, it does however come with challenges. Besides the rare earth materials needed to manufacture photovoltaic cells and wind turbines, there is also a challenge of balancing over-generation during periods of low energy consumption. This claim is exemplified by a 2013 article published by the California Independent System Operator, where the famous "duck-curve" term was introduced.

In this article, concerns were expressed about the ever-increasing level of photovoltaic power generation. In brief, during the spring months, when solar power generation is high, but the surrounding temperature is cold enough not to run air conditioners, there exists a time-frame where over-generation of power occurs. Specifically, this over-generation occurs during midday when solar radiation is at its maximum intensity. In an aggregate sense, the duck-curve shows a decrease in midday net power consumption followed by a steep ramp towards the day's peak power usage as the sun's solar intensity decreases.

As larger amounts of solar energy penetrate the grid, a more significant transitional ramp rate is experienced by the electric utility. One of the more prominently discussed measures to combat the duck curve is wide-scale battery storage. This measure helps improve grid stability by deferring midday solar power generation to hours of peak consumption. The combination of Tesla's Powerwall and integrated solar shingles is one example of residential battery storage. Although great strides have been made in battery storage, it is still somewhat inaccessible to the general public. Furthermore, high infrastructure costs and disruptive storage technologies continue to prevent utilities from harnessing this energy. For clarity, utility companies generate power according to the real-time demand of their end-users. An electric utility's stability is partially linked with the predictability of how much power will be demanded at any given time. Increased battery storage levels help smooth out volatile power demand, reducing the difficulty of predicting power ramp-rate. Suffice to say, batteries increase the flexibility of the grid.

Nevertheless, there are alternative methods of increasing the flexibility of the grid. One such method is the deployment of demand response programs. A demand re-

response program is defined as the systematic limitation of energy consumption by an electric utility to maintain the power generation and distribution equipment within safe operating conditions. Many demand response programs exist, ranging from emergency protocol to adaptive price control, including various ancillary services. The emergency protocol deals with reducing power to prevent blackout/brownout events, where blackouts or rolling brownouts are widespread power outages often caused by the damage of power distribution equipment. Adaptive price control, as will be discussed in Section 2.6, is akin to a soft constraint in an optimization program. By penalizing power consumption during times of peak demand, utility companies are able to smooth out power consumption via the marketplace. Other such demand response events include the tracking of a demand signal produced by an electric utility. Typically, a contractual agreement is made between an electric utility and its participating end-users to manage power consumption by way of limiting certain devices via a home's smart metering device. In return, an electric utility offers rebate on power bills. Both battery integration and demand response programs attempt to accomplish the same task of minimizing peak power consumption and load volatility.

Smart-grid systems are a vast conglomeration of different power generation and distribution related topics. A smart grid system is primarily defined by its ability to balance various energy sources such that the grid's efficiency and resilience is improved. The umbrella of smart-grid systems includes many topics, including the previously mentioned demand response programs and Thermostatically Controlled Load (TCL)s. As the name suggests, TCLs deal with the control of various energy-consuming devices. Because TCLs can be modeled like a battery, many unique control methods have been developed to delay/defer residential appliances' power consump-

tion. In a residential setting, TCLs represent the most significant form of energy consumption (Kwac et al., 2014). Such appliances include water heaters, HVAC systems, and refrigerators, to name a few. These TCL devices create a differential of heat within a medium, usually for purposes of comfort or food storage.

In a black-out scenario, upon reinstatement of power, TCL devices contribute to the grid’s greatest transient effects. The term “transient” is meant to indicate the effect of load synchronization, see Section 4.1.1. Because TCL devices are responsible for consuming the most significant amount of power and have the greatest impact on the grid if powered down improperly, many researchers have focused on controlling said devices. TCL devices store heat, or the lack thereof, in a medium. For instance, air conditioning units remove thermal energy via a compressor to reduce a residential home’s indoor air temperature. When a black-out scenario occurs, a lack of power prevents the removal of heat by the compressor. During this time, the indoor air temperature gradually rises, given a warm outside environment. Upon reinstatement of power, the control algorithm governing this hypothetical HVAC device will begin to cool itself off immediately. If all air conditioning units simultaneously turn on, a resulting large spike in aggregate power will be experienced by the grid. This load synchronization effect is burdensome to utility companies for various reasons. Often this load synchronization is challenging to predict and even more so challenging to meet the demand. Likewise, these effects are observed in other TCL devices. Preventing blackouts is a high priority for electric utility companies due to the potential damage to power transmission equipment.

In congruence with the past century’s computational advancements, various control methodologies have been proposed for TCLs. It is helpful to further subdivide

these control methodologies by how they are implemented and what they attempt to do. One such classification looks at how these devices are controlled, for instance, a top-down approach from an electric utility or the emergent characteristics of a decentralized control framework. With regard to a **centralized controller**, a utility samples a population of TCL, computes their control actions, and distributes the command. **Decentralized control** frameworks tackle the control of TCL differently. Each TCL agent is responsible for computing its own control action. This decentralized approach is further augmented with localized communication, as will be shown in Section 2.2.

Much of the literature in TCLs stems from the seminal work produced by (Malhame & Chong, 1985). In this paper, Malhamé and Chong use a Markovian hybrid-state model to describe the switching dynamics of a population of homogeneous electric space heaters. When aggregated, a system of Fokker-Planck equations arises, whose solution allows for the statistical aggregation of TCL devices. This work served as the foundation for many other papers. Two decades later, (Callaway, 2009) furthered Malhamé and Chong’s work by deriving an exact solution to the heterogeneous hybrid state aggregated model (Fokker Planck equations). In Callaway’s hybrid state model, a first-order ETP model is used to describe a TCL device’s temperature evolution. (Zhang et al., 2013) further improved the accuracy of Callaway’s hybrid-state framework by modeling a residential system’s air conditioning unit as a second-order ETP model. This updated model describes the coupled relation between a home’s indoor air and bulk mass temperatures, thereby improving the predictive accuracy of the model. In addition, Zhang et al. took a unique approach to the aggregation of air conditioning units by using a state bin transition framework to track an aggregate

demand reference signal while maintaining indoor air temperatures at/near their appropriate settings. Additional state bin transition frameworks are also explored in the paper produced by (Koch et al., 2011) and the report produced by (Liu & Shi, 2015). These state bin transition methods control the aggregate power demand by tracking a reference signal generated by a utility company, where aggregate demand is a function of each TCL device’s power consumption.

Another prominent area of research in the field of smart grid systems is adaptive price control. Techniques such as the ones proposed by (Zhou et al., 2017) introduced a control framework that, on a two-part basis, minimized energy and capacity costs of an aggregate power signal via model predictive control while using a temperature priority list to control each TCL device. Another adaptive price control strategy, developed by (Behboodi et al., 2018), used an agent-based modeling approach to induce a demand response event. Behboodi used a transactive control paradigm to regulate when TCLs consume power. (Wang et al., 2019) proposed a similar control architecture based on the dynamic cost of power. Adaptive price control and agent-base modeling continues to be active fields of research.

Several papers have proposed the use of reinforcement learning to learn the complex state-action space of TCL. Reinforcement learning, a branch of machine learning, deals with how independent actors interact within their environment. These actions either maximize a reward, or conversely minimize a penalty, which is defined by its objective function. Relative to other control architecture, machine learning and, by extension, reinforcement learning is a relatively new field of study, with many papers being published at the time of this document. Some of the more prominent articles in the field of TCLs and reinforcement learning are (Ruelens et al., 2016), (Ruelens

et al., 2019), and (Kazmi et al., 2019). A major advantage of reinforcement learning is that no predictive model is needed. However one of the main difficulties is constructing an objective such that the model learns appropriate control actions.

The underlining motivation of a decentralized controller is the reduced computational complexity and respect for end-user privacy, see Section 2.2. As discussed above, decentralized control requires individual agents to compute their control actions such that the systems’ constraints are satisfied. As it relates to decentralized TCLs, the work of (Tindemans et al., 2015) describes a method of individual TCLs tracking relative power reference signals. Although similar in scope to this research, a critical difference is that this paper does not require the tracking of an aggregate reference signal to achieve reductions in aggregate power consumption. As will be discussed in Section 3, minimization of aggregate power and its associated volatility is partially accomplished via Peer-to-Peer (P2P) communication. The work of (Liu & Shi, 2015) also provides a decentralized control framework for communicating HVAC systems. In this work, a building complex’s HVAC system is tasked with optimally cooling each individual tenant space. Liu used P2P communication to control the rate of cooling in each respective unit of the building complex. As opposed to the approach used in Liu, our method takes a more holistic approach by simulating residential homes. Furthermore, the optimization programs used to determine an HVAC’s control effort is formulated in a fundamentally different way.

In (Kuwada et al., 2020) a modelless decentralized control framework is presented, which uses sparse P2P communication. More accurately, an advanced hysteresis controller is presented, which statistically distributes the control action of each residential home such that load aggregation is reduced. Similar to Kuwada’s work, this research

aims to use sparse network communication but instead opts to use Model Predictive Control (MPC) to optimally schedule the control actions of a population of TCLs. Due to its similarity, the advanced hysteresis controller proposed by Kuwada will also be introduced in this work and will be used as a comparative metric.

The rest of this thesis is organized as follows, in Section 2, an equivalent thermal parameter model is presented for both SSHP and VSHP. Based on their ETP model, a recursive least squares algorithm is presented to update these characteristic parameters adaptively. A simulated demand response event is introduced, and various control algorithms, like the advanced hysteresis controllers described by Kuwada and Proportional Integral Derivative (PID) controllers, are defined. Lastly, a structure of communicating TCL is developed. In Section 3, four optimal control frameworks are developed for the previously mentioned SSHP and VSHP air conditioning devices. For each TCL device, both a decentralized and centralized network structure is implemented. Using the aforementioned lower-level control frameworks defined in Section 2.3, a baseline study is conducted. Thereafter, all four optimal control frameworks of Section 3 are simulated. Lastly, concluding remarks are given about each controller's feasibility and performance traits followed by future research pathways.

CHAPTER 2:

BACKGROUND

According to the U.S. Energy Information Administration, a majority of the grid's generated electricity is consumed by TCLs. TCLs resemble leaky batteries in that they are characterized by their ability to store thermal energy. For instance, water heaters heat incoming water in an insulated storage tank within a prescribed temperature range where it awaits its use. In addition to water-heaters, other prominent high energy-consuming TCLs are HVAC systems and refrigerator units. It should be noted these TCL devices, including many others, operate on a similar premise. Typically, TCLs are governed by simplistic toggle conditions, like an on/off controller, also referred to as a hysteresis or relay control. In cases where a TCL device can apply continuous control, a PID architecture is used.

Because thermal energy is stored within a medium, there exists a level of flexibility when such energy is consumed/replenished. An electric utility can achieve certain beneficial aggregate characteristics through the proper coordinated control of such devices, as demonstrated in this work. For instance, the optimal control framework of Section 3 shows promising results in reducing aggregate power consumption and its associated load volatility.

In this study, four unique control algorithms are developed for two TCL devices. These TCL devices are single-stage heat pumps, and variable speed heat pumps, com-

monly referred to as SSHP and VSHP respectively. SSHPs and VSHPs are types of HVAC units that differ in the way energy is applied into the system. The key difference between these two devices is that VSHPs can modulate their cooling capacity, whereas SSHPs are either ON or OFF. The main reason for choosing these two TCL devices is that in the United States SSHPs and VSHPs constitute the majority of HVAC systems in residential use. In Section 3, a mathematical model is developed for both of these air conditioning types. These models provide a causal relationship between the control actions of a population of air conditioning devices and the aggregate power consumption experienced by an electric utility. Although this research primarily focuses on SSHP and VSHPs, the mathematical principles described can be further expanded to other similarly controlled TCL devices, like water heaters and refrigerators.

2.1 Modeling Approach

Each control framework for both SSHPs and VSHPs begin by describing a typical residential home's thermal model. In this research, we opt to use a second-order equivalent thermal parameter model (ETP) to describe the temperature dynamics of a residential home following (Zhang et al., 2013). Utilizing the thermal circuit shown in Figure 2.1, two coupled first-order differential equations are formed.

$$C_A \dot{T}_A = Q_A - \frac{1}{R_1} (T_A - T_o) - \frac{1}{R_2} (T_A - T_M), \quad (2.1)$$

and

$$C_M \dot{T}_M = Q_M + \frac{1}{R_2} (T_A - T_M), \quad (2.2)$$

Figure 2.1 Second-order equivalent thermal parameter circuit diagram.

where T_A and T_M denote the interior air and lumped mass temperatures of a residential home respectfully. Adjacent to the residential control volume is the surrounding outside temperature, T_o . Both the outdoor temperature, T_o , and the internal heat-generating elements, denoted Q_M , represent disturbances to this thermal system. Rejection of such disturbances is accomplished via the home's HVAC system, denoted Q_A . The term Q_A symbolically represents a home's HVAC system's cooling capability. In particular, Q_A represents the rate of heat removed from the interior air. The elegance of this ETP model is that four measurable parameters fully define it, that being the thermal capacities C_A , C_M and the thermal resistance properties R_1 , R_2 . These parameters map to measurable characteristics of the residential system.

Based on these thermal parameters, this ETP model describes the coupled relationship between the fast-changing indoor air temperature, T_A , and the slower bulk mass temperature, T_M , of a home. This second-order ETP model also provides a unique method of controlling such a device. Because energy is stored within the residential system's interior space, this ETP model acts as a battery. The control effort, denoted m , for both SSHP and VSHP devices is determined based on the amount of thermal energy captured within the system. Because this system operates on relatively slow time-scales, minutes to hours, the amount of cooling can be controlled so that temperature deviations are unnoticed by the end-user. Additionally, unlike the less accurate first-order ETP model, this model framework provides a valid account of how energy traverses throughout the residential system, further improving the predictive capability of the control frameworks presented in Section 3.

Following (Fuller, 2017), Equations (2.1) and (2.2) are manipulated to form a

second-order differential equation in terms of T_A and its derivatives $\dot{T}_A$ and $\ddot{T}_A$,

$$\begin{aligned} C_M C_A R_2 \ddot{T}_A + \left(C_M \left(\frac{R_2}{R_1} + 1 \right) + C_A \right) \dot{T}_A + \frac{1}{R_1} T_A \\ = \frac{C_M R_2}{R_1} \dot{T}_o + \frac{1}{R_1} T_o + C_M R_2 \eta \dot{m} + \eta m. \end{aligned} \quad (2.3)$$

The complete derivation of this second-order ETP model, using Equations (2.1) and (2.2), is provided in Appendix A.1. In Section 2.4, a Recursive Least Squares (RLS) algorithm is introduced to adaptively estimate the parameters defining this second-order ETP model. Before the RLS algorithm is introduced, the constant terms of Equation (2.3) are further abbreviated by elements of a parameter vector, $\theta \in \mathbb{R}^5$. This redefinition of Equation (2.3) has the form,

$$\theta_1 \ddot{T}_A + \theta_2 \dot{T}_A + \theta_3 T_A = \theta_4 \dot{T}_o + \theta_5 T_o + \theta_5 \eta \dot{m} + \eta m. \quad (2.4)$$

Within Equations (2.3) and (2.4), the terms Q_A and $\dot{Q}_A$ are replaced by ηm and $\eta \dot{m}$ respectively. The term η represents the heat removal capacity of a home's HVAC system and m , a controllable parameter, scales η according to the governing control algorithm. The term Q_M , similar to that of Q_A , represents alternative heat sources/sinks. In a typical residential setting, Q_M includes, but is not limited to, the heat generated by home appliances, solar radiation, and the conductive heat transfer between the home's lumped mass and ground. Due to the minimal effects of Q_M , especially concerning $Q_A := \eta m$, its inclusion will be neglected in both SSHP and VSHP models. Lastly, the outside temperature, T_o , and its derivative, $\dot{T}_o$, represent a disturbance to the thermal system. Rejection of such disturbances is paramount to a successful control algorithm. In Sections 2.1.1 and 2.1.2, a state-space representation

Figure 2.2 SSHP, $m \in \{0, 1\}$, and VSHP, $m \in [0, 1]$, control states, where the SSHP is either Off or ON and VSHP is linearly between OFF and ON.

of Equation (2.4) is developed for both SSHPs and VSHPs.

2.1.1 SSHP Dynamic Model

As shown in Figure 2.2, a SSHP is a type of HVAC unit that toggles OFF or ON depending on the indoor air temperature, T_A , relative to its upper and lower dead-band thresholds, respectfully denoted δ^+ and δ^- . This toggle condition is mathematically defined as,

$$m(t) = \begin{cases} 0, & \text{if } T_A \leq \delta^- \\ 1, & \text{if } T_A \geq \delta^+ \\ m(t - \epsilon), & \text{otherwise,} \end{cases} \quad (2.5)$$

where the dead-band thresholds are further defined as $\delta^+ = T_{SP} + \frac{\Delta T}{2}$ and $\delta^- = T_{SP} - \frac{\Delta T}{2}$. An end-user determines the temperature set-point, T_{SP} , of the home while

the dead-band width, ΔT , not to be confused with the time-step length, Δt , is set by the control algorithm. As previously stated, the term m is the control input for a SSHP device. Moreover this control term is binary, i.e., can only assume a value of either 0 or 1. While ON, the SSHP device removes heat from the interior air at the cooling rate, η , otherwise the system is OFF. Due to the binary nature of the control input, $m \in \{0, 1\}$, the ETP model of Equation (2.4) can be simplified. In this scenario, the control input's rate of change, $\dot{m}$, consists of a series impulses as Equation (2.5) toggles m OFF and ON. We found the term $\dot{m}$ to have a negligible effect on the overall system during simulation. Moreover, this term was also neglected in the various literature results surveyed, see (Fuller, 2017). By removing $\dot{m}$ from Equation (2.4), the new ETP model dynamics for SSHPs has the form,

$$\theta_1 \ddot{T}_A + \theta_2 \dot{T}_A + \theta_3 T_A = \theta_4 \dot{T}_o + \theta_3 T_o + \eta m, \quad (2.6)$$

where the new parameter vector for SSHPs is $\theta \in \mathbb{R}^4$. The state space representation of Equation (2.6) has the form,

$$\dot{\mathbf{x}} = A\mathbf{x} + B \left(\eta m + \theta_4 \dot{T}_o + \theta_3 T_o \right), \quad (2.7)$$

where $\mathbf{x} = \begin{bmatrix} T_A & \dot{T}_A \end{bmatrix}^\top$ and,

$$A = \begin{bmatrix} 0 & 1 \\ -\frac{\theta_3}{\theta_1} & -\frac{\theta_2}{\theta_1} \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ \frac{1}{\theta_1} \end{bmatrix}.$$

In order to use the model dynamics of Equation (2.7) in the controller proposed

in Chapter 3, this state space representation must first be discretely written using forward Euler method.

$$\mathbf{x}^{k+1} = (I + \Delta t A) \mathbf{x}^k + \Delta t B \left(\eta m^k + \theta_4 \dot{T}_o^k + \theta_3 T_o^k \right), \quad (2.8)$$

where k represents the simulation's current time-step and Δt similarly represents the elapsed time between subsequent time-steps of integration.

2.1.2 VSHP Dynamic Model

As shown in Figure 2.2, VSHPs are capable of providing continuous control values between 0 and 1. The amount of heat removed from the indoor control volume is similarly ηm , however, now the cooling capacity is linearly scaled. A VSHP and the encompassing home are modeled with Equation (2.4). Unlike the SSHP's ETP model, the VSHP ETP does not neglect $\dot{m}$.

As will be further shown in Chapter 3, we choose to use $\dot{m}$ as the control variable and augment the state with m . Managing a VSHP using $\dot{m}$ allows the control algorithm to reduce abrupt changes in control effort when properly minimized. Our method of control is novel and shows promising results in the regulation of indoor air temperatures.

The state space representation of Equation (2.4) has the form,

$$\begin{aligned} \dot{\mathbf{x}} &= A\mathbf{x} + B \left(\eta(m + \theta_5 \sigma) + \theta_4 \dot{T}_o + \theta_3 T_o \right), \\ \sigma &= \dot{m}, \end{aligned} \quad (2.9)$$

where $\mathbf{x} = \begin{bmatrix} T_A & \dot{T}_A \end{bmatrix}^\top$ and,

$$A = \begin{bmatrix} 0 & 1 \\ -\frac{\theta_3}{\theta_1} & -\frac{\theta_2}{\theta_1} \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ \frac{1}{\theta_1} \end{bmatrix}.$$

Through the application of forward Euler method, Equation (2.9) is thereby rewritten in a discrete format as,

$$\begin{aligned} \mathbf{x}^{k+1} &= (I + \Delta t A) \mathbf{x}^k + \Delta t B \left(\eta (m^k + \theta_5 \sigma^k) + \theta_4 \dot{T}_o^k + \theta_3 T_o^k \right), \\ m^{k+1} &= \sigma^k \Delta t + m^k. \end{aligned} \quad (2.10)$$

Again, Δt and k denote the step-length and simulation time-step respectively.

2.2 Network Communication

Much of the proposed literature in smart grid systems, particularly TCLs, are constructed using a centralized framework. In such a framework, a utility service provider determines the control action of a population of TCLs, typically through tracking an aggregate reference signal or adaptive price control. This methodology, including various other centralized control frameworks, has shown good performance benefits. However, there are several glaring drawbacks to a centralized framework. Some of the more prominent challenges include its scalability, vulnerability to cyber-attack, and inherent lack of consumer privacy. These challenges are addressed by using a decentralized control framework. A decentralized framework relies on the autonomy of the participating agent to calculate their control action. As will be demonstrated in Chapter 4, the performance of a decentralized controller is further improved through

the localized communication between neighboring TCLs.

Before the network communication model is presented, we will first justify its structure. As discussed above, a major challenge for centralized controllers is the lack of consumer privacy. This lack of consumer privacy comes from the data's availability, which might be obtained through nefarious means or given the right market conditions distributed by the electric utility itself. We address these privacy concerns by regulating the number of individual communication pathways between neighboring homes and consider the proximity in which connections are made. Communication is limited to what might be ascertained if a residential homeowner were to open their window and listen when a neighboring air compressor turns OFF or ON. Alternatively, in the case of VSHPs, one can listen to the amount of cooling being applied. This communication style addresses the privacy concerns of its end-users and inadvertently reduces exposure to a cyber-attack. Lastly, the information shared between neighboring TCL devices only relates to the control input, m , not other state variables, T_A and $\dot{T}_A$.

By deduction, this communication network implies that the control effort of the p^{th} agent, m_p , is partially influenced by the control action of its neighboring homes, $m_i \forall i \in \mathcal{N}_p$, where $\mathcal{N}_p$ represents the set of homes connected to the p^{th} agent. As will further become helpful later, the population set is similarly defined as $\mathcal{N} = \{1, \dots, N\}$ and represents all of the HVAC systems within the utility companies' purview, $\mathcal{N}_p \subseteq \mathcal{N}$.

The communication structure described above can be modeled using the random regular graph structure of Figure 2.3. In the literature of graph theory, other graph structures exist and are used to model specific phenomena. For instance, a star graph

Figure 2.3 Unweighted random regular graph ($N = 50$, $N_{cd} = 4$)

loosely resembles that of a centralized controller. Furthermore, graph structures like the scale-free graph proposed by (Barabási & Bonabeau, 2003) is widely applicable to natural phenomena like neuron mapping and city structure. As it relates to this study, a random regular graph with a finite number of links between TCLs most accurately depicts the limited communication posed by our control framework. Similar to Figure 2.3, the connection density, N_{CD} , of each simulated graph is limited to four connections, i.e., the magnitude of the neighbor-set $|\mathcal{N}_p| = 4 \forall p$. The effects of these

parameters are subjects of future avenues of research.

For context, graphs are commonly presented either as an ordered pair of vertices and edges, $G = (\mathcal{N}, E)$, or as elements of an adjacency matrix, $\mathcal{A}$. For the former, vertices, denoted $\mathcal{N}$, are the set of all nodes/points within the structure, and edges, denoted E , are defined as $E \subseteq \{\{a, b\} \mid a, b \in \mathcal{N} \wedge a \neq b\}$. Alternatively, in the latter case, an adjacency matrix, denoted $\mathcal{A}$, is a square matrix whose elements depict a weighted edge between two vertices. As it relates to simple undirected graphs, the weighted elements are either 0 or 1, where 0 denotes no connection and 1 represents a connection. The p^{th} row of the adjacency matrix, $\mathcal{A} \in \mathbb{R}^{N \times N}$ represents the neighbor-set $\mathcal{N}_p$. Additionally, the connection density for the p^{th} TCL is defined as $N_{CD} = \sum_{j=1}^N \mathcal{A}_{p,j}$.

The communication described in this research is the transfer of bitwise control input data along edges. This communication necessarily implies that a pipeline exists between TCLs and is capable of data transfer. Several communication pipelines exist, such as direct line communication and services provided via Internet-of-Things. The emergence of smart thermostats is one such example.

2.3 Baseline Control Frameworks

Similar to the optimal control frameworks presented in Section 3, the advance hysteresis controllers presented in (Kuwada et al., 2019, 2020) is built using the network structure described in Section 2.2. This advanced hysteresis control framework is used as a comparative metric against the optimal SSHP controllers defined in Section 3. The advanced hysteresis controller begins by defining the term $\tilde{m}$ to represent the average control effort among the p^{th} agent's neighbor-set $\mathcal{N}_p$. This is done systemat-

ically with,

$$\tilde{\mathbf{m}}^k = \frac{1}{N_{CD}} \mathcal{A} \mathbf{m}^{k-1}. \quad (2.11)$$

In Equation (2.11), $\mathbf{m} \in \mathbb{R}^N$ is a vector of the current control states of the population of SSHPs. Using $\tilde{m}$ calculated with Equation (2.11), the advanced hysteresis controller is mathematically described as,

$$m_p^k = \begin{cases} 0, & \text{if } \psi_p^k - K_g \tilde{m}_p^k \leq 0 \\ 1, & \text{if } \psi_p^k - K_g \tilde{m}_p^k \geq 0 \\ m_p^{k-1}, & \text{otherwise,} \end{cases} \quad (2.12)$$

where K_g is a proportional constant term and,

$$\psi_p^k = \frac{T_{A,p}^k - \delta_p^-}{\delta_p^+ - \delta_p^-}, \quad (2.13)$$

is the normalized indoor air temperature relative to its dead-band values, δ^+ and δ^- . The advanced hysteresis controller of Equation (2.12), in essence, adaptively adjusts each TCL's dead-band width, ΔT , proportionally to the number of ON states of the p^{th} agent's neighbor-sets, $\mathcal{N}_p$. More time will be spent idly warming up by expanding a TCL's dead-band width. When controlled by the hysteresis controller of Equation (2.5), there exist conditions where load synchronization may occur. The advanced hysteresis controller combats load synchronization by stochastically defusing control actions of the population set $\mathcal{N}$. The reader might ask "why not stop with this advanced controller?". During simulation, load synchronization was still observed in the system. Moreover, this diffusion method trades end-user comfort for a reduction

Figure 2.4 Closed loop controller for VSHP.

in load synchronization.

The decentralized controllers proposed in Section 3 attempts to find an optimal solution that balances aggregate power characteristics with end-user comfort. Simulation results are provided in Section 4 for the two SSHP lower-level controllers, Equations (2.5) and (2.12), and a PID control framework for VSHPs defined next.

In addition to the discrete control frameworks described by Equations (2.5) and (2.12), a continuous control framework is included to compare the simulated results of the VSHP controller. Relating to real applications, PID controllers overwhelmingly represent the market-space. Due to its model-free design, relative ease to incorporation, and general high accuracy PID integration is an attractive option for continuous control scenarios. In simplistic terms, PID control reduces the error between the controller's current value and the desired position of where it's supposed to be. For instance, in temperature control, an end-user sets the thermostat's reference temperature, T_{SP} . Based on the system's actual temperature, heat is added/removed such that these two values coincide.

The closed loop system observed in Figure 2.4 uses the error, defined as the

difference between the set-point and indoor air temperature,

$$e(t) = T_{SP} - T_A, \quad (2.14)$$

to calculate the control action, $u(t)$. In the time-domain, this control action $u(t)$ is defined as,

$$u(t) = K_P e(t) + K_I \int_0^t e(\tau) d\tau + K_D \frac{d}{dt} e(t). \quad (2.15)$$

Each cumulative term controls a particular attribute of the response. The first term, $K_P e(t)$ increases the control response proportionally to the error $e(t)$. Next, the integral term, $K_I \int_0^t e(\tau) d\tau$, regulates control actions according to the error built up in the system. Lastly, the derivative term, $K_D \frac{d}{dt} e(t)$, attenuate the error's rate of change. These three PID terms allow for the tracking of the set-point temperature and represents what is most likely used to control modern day VSHPs. Before the control action is sent to the plant, its value must first be bounded by the natural saturation limits of the VSHP device. Both SSHPs and VSHPs are limited to the maximum cooling capacity η corresponding to $m = 1$. The saturation block limits $u(t)$ between 0 and 1 where,

$$m(t) = \max(0, \min(u(t), 1)) \quad (2.16)$$

This saturation constraint is duly implemented as a clamp function. After calculating $m(t)$, the control action is sent to the VSHP device, also referred to as the plant, where it is then acted upon. Integrated sensors measure the temperature of the internal control volume. The measured value is then used in the following iteration

to calculate the next error signal, thereby closing the PID loop.

The reader might ask “why would a model-based optimal controller be necessary given the wide success of PID systems”. Like the hysteresis controllers described above, this PID controller is prone to load synchronization under certain simulated conditions. Moreover, such a controller fails to account for neighboring control action. This results in suboptimal aggregate power consumption performance.

2.4 Parameter Estimation

A major challenge associated with any model-based control is the proper modeling of the plants’ dynamical characteristics. In this study, the non-linear temperature dynamics of a residential home are reasonably approximated with second-order ETP models, shown by Equations (2.8) and (2.10). Before the optimal control of both SSHP and VSHP devices, measurements of the thermal parameters defining the ETP models must be made. Whether by inference or randomly assigned, some degree of measurement error is always assumed to exist. When poorly estimated, the predictive model, which integrates state values, will result in suboptimal controller performance. Our parameter estimation algorithm aims to nullify the initial measurement error and to track long-term change caused by parameter drift. As previously discussed, by replacing each term of the ETP model with elements of a parameter vector θ , the newly formed linear model then becomes a prime candidate for the recursive least squares algorithm RLS see (Åström & Wittenmark, 2013).

As will be further discussed in Chapter 3, our proposed optimal control frameworks are inherently more complex when compared to the lower-level controllers of Section 2.3. Due to this increased computational burden, more time is required to solve for the control actions of a TCL. This RLS approach is a closed-form solution to a classic

minimization of the sum of squares problem. The recursive least squares algorithm modifies system parameters such that the difference between the plant's output and the predicted response of the model is minimized. In conjunction with the novelty of the control approach for VSHPs, incorporating a parameter estimation model in this line of research is deemed unique. As shown in Sections 2.1.1 and 2.1.2 two ETP models are presented therefore two RLS algorithms are needed. More specifically, an RLS algorithm will be developed for both SSHP and VSHP unit type. This RLS algorithm begins by re-defining Equations (2.6) and (2.3) in terms of regressor and observer variables. The notation presented matches closely (Åström & Wittenmark, 2013); however, some modifications are made.

$$y_k = \varphi_k^\top \theta_0, \quad (2.17)$$

where the observer, y_k , and regressor, φ_k , terms for the SSHP and VSHP models are respectively defined as,

$$y_k = \eta m^k, \quad (2.18)$$

$$\varphi_k^\top = \begin{bmatrix} \ddot{T}_A^k & \dot{T}_A^k & (T_A^k - T_o^k) & -\dot{T}_o^k \end{bmatrix},$$

and,

$$y_k = \eta m^k, \quad (2.19)$$

$$\varphi_k^\top = \begin{bmatrix} \ddot{T}_A^k & \dot{T}_A^k & (T_A^k - T_o^k) & -\dot{T}_o^k & -\eta \sigma^k \end{bmatrix}.$$

Notice the VSHP's reformulation includes $-\eta \sigma^k$ where σ^k represents the rate at which the control input, m^k , changes. The term ηm^k is chosen to be the observer variable.

Whether physically measured, empirically determined, or in the case of this study,

generated about a known statistical distribution, all elements of the thermal parameter vector, θ , must be known prior to running the RLS algorithm. This initial thermal parameter vector is denoted θ_0 and is further defined as,

$$\theta_0 = \begin{bmatrix} \theta_1^0 & \theta_2^0 & \theta_3^0 & \theta_4^0 \end{bmatrix}^\top \in \mathbb{R}^4, \quad (2.20)$$

and

$$\theta_0 = \begin{bmatrix} \theta_1^0 & \theta_2^0 & \theta_3^0 & \theta_4^0 & \theta_5^0 \end{bmatrix}^\top \in \mathbb{R}^5. \quad (2.21)$$

Before the RLS algorithm can predict thermal parameters with sufficient accuracy, both the observer and regressor terms are to be collected over a set of initial time-steps $k \in \{1, \dots, k_s\}$ to form the following observer and regressor matrices,

$$Y_{k_s} = \begin{bmatrix} y_1 \\ \vdots \\ y_{k_s} \end{bmatrix}, \quad \Phi_{k_s} = \begin{bmatrix} \varphi_1^\top \\ \vdots \\ \varphi_{k_s}^\top \end{bmatrix}. \quad (2.22)$$

Relegating the number of data points needed prior to running the RLS algorithm ensures enough information is captured within Y_{k_s} and Φ_{k_s} to faithfully predict new parameter values. The last time-step, $k_s \in \mathbb{N}$, is chosen such that $\Phi_{k_s}^\top \Phi_{k_s}$ is non-singular. Given the initial starting value, θ_0 , and adequately populated observer and regressor matrices, used to initialize the matrix, $P_{k_s} = (\Phi_{k_s}^\top \Phi_{k_s})^{-1}$, the RLS algorithm has the following sequence of events,

Algorithm 1 RLS with Exponential Forgetting

- 1: Initialize $P_{k_s} = (\Phi_{k_s}^\top \Phi_{k_s})^{-1}$
 - 2: **while** $k_s < k \leq K$ **do**
 - 3: $S_k = P_{k-1} \varphi_k (\lambda + \varphi_k^\top P_{k-1} \varphi_k)^{-1}$
 - 4: $P_k = (I - S_k \varphi_k^\top) P_{k-1} / \lambda$
 - 5: $\theta_k = \theta_{k-1} + S_k (y_k - \varphi_k^\top \theta_{k-1})$.
 - 6: **end while**
-

Observing Algorithm (1), the RLS sequence starts by initializing the properly populated matrix, $P_{k_s} = (\Phi_{k_s}^\top \Phi_{k_s})^{-1}$. Necessarily, the governing controller must collect the observer and regressor term for k_s iterations before updating θ_0 . A circular dependency exists between k_s and the matrix P^k . A logical solution to combat this circular dependency is to check the rank of P^{k_i} for each time-step $k_i \in \{0, \dots, k_s\}$. Once rank sufficient, the current time-step, k_s , is the point when the RLS algorithm begins. Paraphrasing (Åström & Wittenmark, 2013), the RLS algorithm communicatively updates the previous time-steps thermal parameters, θ_{k-1} with $S_k(y_k - \varphi_k^\top \theta_{k-1})$. After being updated, the newly estimated thermal parameter vector, θ_k , more accurately represents the observer and regressor data used.

By deduction, thermal parameters will converge upon the theoretical values of the plant. In our simulation, the plant response is also simulated; however, in reality, the system's actual dynamics is expected to be non-linear or at the very least capable of change. An exponential forgetting variant of the RLS algorithm is presented to combat non-linear effects and the change due to degradation or renovation. This RLS with exponential forgetting, defined by $\lambda \in [0, 1]$, allows the RLS algorithm to weigh previous data's effect. For reference, while λ approaches 1 the RLS algorithm with

exponential forgetting becomes the vanilla RLS algorithm.

During experimentation, this RLS algorithm was found to produce high-frequency noise. A low-pass filter was used to combat both the general system and measurement noise. More explicitly, newly estimated thermal parameter values, θ_k , were altered with a low-pass filter to attenuate high frequency noise injected from system and measurement level noise. The low-pass filter used during simulation is defined as,

$$\theta_k = a_0\theta_k + \sum_{i=1}^4 b_i\theta_{k-i}, \quad (2.23)$$

where the low-pass filter constants are respectfully defined as,

$$C_{LPF} = \begin{bmatrix} a_o \\ b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix} = \begin{bmatrix} (1-x)^4 \\ 4x \\ -6x^2 \\ 4x^3 \\ -x^4 \end{bmatrix},$$

and,

$$x = e^{-14.445f_c}.$$

Based on the signal generated by newly estimated thermal parameters, the frequency cutoff was empirically set as,

$$f_c = \frac{\Delta t}{3}.$$

2.5 SSHP Lockout Procedure

In this research, SSHP devices are kept within safe operating conditions by limiting the effects of short cycling of the compressor. An HVAC system removes heat by passing warmer air over an array of colder evaporative coils arranged in such a manner that efficiently exchanges heat. As warm air passes over these evaporative coils, captive refrigerant heats up, thereby inducing a phase change. The latent heat removed from this air causes the refrigerant to evaporate/boil. Due to the characteristics of fluids undergoing a phase change, a large amount of thermal energy is removed from the system. Once evaporated, this warm gaseous refrigerant is routed outside towards the fan. In most residential settings, this rather large fan is located on the periphery of a home. Within this air conditioning unit is a compressor. The purpose of this compressor is two-fold. Firstly, this compressor moves refrigerant throughout the evaporator and condenser lines by inducing a pressure differential. Secondly, when the warm evaporated refrigerant is compressed, so too is the captive thermal energy. This region of high-temperature fluids is then passed through the condensing coils, where a large fan removes this heat, thereby ejecting it into the ambient outside air. Because the refrigerant is pressurized at a temperature conducive to a phase change, the refrigerant is thereby turned back into a liquid state. This cycle then repeats until an optimal indoor air temperature is achieved. Analogous to how the compressor pumps fluid, the SSHP system, as a whole, pumps concentrated indoor heat outside. Much like a beating heart takes oxygen to the brain, the compressor within the HVAC system plays a vital role in this heat removal process. When an HVAC experiences short-cycling, the compressor's operational lifespan is greatly diminished. Short-cycling is a process where the compressor is turned OFF,

$m = 0$, and ON, $m = 1$ in quick succession. The rate at which the compressor turns OFF and ON leads to mechanical wear.

The phenomenon of short-cycling is an important reason why the cooling capacity, η , must be tailored to the thermodynamic properties of a home. Short-cycling occurs when the HVAC system's rated tonnage is much greater than the heat that can reasonably be removed from the system. Increasing the rate at which heat is removed thereby reduces the time needed to cool, given all other properties of the home remain constant. Undersized HVAC units fall victim to the same mechanical damage too. However, unlike the rapid toggling of control states, an under-sized HVAC unit struggles to maintain its set-point temperature under harsh ambient outside conditions. By deduction, we postulate that there exists an optimally sized HVAC cooling capacity, η , for a given residential home. Typically, HVAC units are sold in half-ton increments, where one ton of cooling is the amount of heat needed to melt two-thousand pounds of water in a twenty-four-hour period.

In our simulation, the discrete cooling capacity is chosen in such a manner that cools a home from its upper dead-band, δ^+ , to its lower dead-band, δ^- , in a thirty-minute time frame. We prevent the phenomenon of short-cycling by subjecting all SSHP devices to a lockout period, Δt_l . During the lockout period, the control effort must remain OFF, beginning when it toggles from OFF to ON. This lockout period is implemented as an optimization constraint in Section 3. Before a lockout scenario, a logic check identifies if the MPC controller has toggled. If so, for a period, Δt_l , the control action, m , is prevented from switching back ON. Given an outside temperature, T_o , and an oversized cooling capacity, η , which causes short cycling, a lockout constraint like the one described above limits the overall toggle frequency. In

doing so, the mechanical wear of the compressor is minimized. The lockout period is capable of being adjusted according to the manufacturers' specifications.

2.6 Demand Response

Amongst many other industry challenges, an electric utility has the difficult task of employing demand response events, also referred to as conservation events. The electric utility brings forth these events to incentivize their customers to reduce power consumption during the day's peak hours. Before our specialized demand response event is discussed, it is helpful to understand how the electric grid is structured. Power generation equipment can be classified into three main categories. These categories are the base, intermediary, and peak loads.

The base-load supplies most of the power consumed by the grid. Typically, the base-load is composed of power generating equipment that slowly reacts to change, for instance, large Rankine cycle, nuclear fission reactors, and hydroelectric plants. These power plants make up the entirety of the non-cyclical section of an aggregate demand signal. They are often characterized by the amount of power supplied before the intermediary-load must be tracked. Next is the aggregate power consumption deemed to be cyclical in nature. This intermediary-load is powered by less efficient power generating equipment. Still, it reacts much faster to changes in the aggregate demand signal. The intermediary-load comprises equipment such as diesel power generators. These power sources are responsible for tracking the majority of the intermediary aggregate demand signal. Lastly and arguably the most challenging load to predict is peak power demand. This exuberant demand is characterized by its fast reaction to power demand swings.

The purpose of peak power generation is to provide power not only during high

power consumption but also to pick up the slack during high load volatility. The power generation equipment responsible for supplying peak power is the least efficient of the three load types. Because peak power generators pick up the slack for highly volatile and complex to predict regions of the aggregate power consumption signal, they are by deduction most at risk of failure.

It can now be understood why an electric utility desires a smooth aggregate demand signal. Firstly, it allows a more significant portion of power to be produced by the base-load, which is less expensive to produce. Conversely, less energy is being produced by the more costly intermediary and peak power generating equipment. Secondly, if the aggregate demand is smoothed, there is less risk of mechanical damage to these adaptive power plants, thereby improving the electric grid's resilience and reliance. Thirdly, tracking aggregate power consumption becomes a more manageable problem when gradual transitions in the aggregate signal are experienced. The electric utility may then employ more sophisticated control problems to balance these three plant types. For these reasons, it is highly desirable and lucrative for an electric utility that TCLs be optimally controlled such that the aggregate power consumption is predictably smooth.

Now that an electric grid's power generating equipment has been explained, the question is "what might an electric utility do to mitigate these troublesome intermediary and peak loads?". A market-based solution is that an electric utility incentivize its customers not to use power during the day's peak hours. Usually, an electric utility will offer rebates for the contractual right to either shift or limit its participating clients' power consumption. Many forms of demand response events have been proposed in literature ranging from adaptive price control to emergency protocol.

This study focuses on emergency demand response events. Reasons for an emergency demand response event are that an electric utility anticipates higher power consumption than what can be reasonably produced. An electric utility limits specific devices' power consumption through the customers' smart meter device in this emergency event. Some of the more easily controlled devices are HVAC units and water heaters.

As previously discussed, HVAC systems store thermal energy within the residential control volume. An electric utility can metaphorically tap into this energy reservoir by preventing the HVAC system from running, thereby limiting troublesome peak demand. If done correctly, a reduction in peak demand is traded for mild end-user discomfort, i.e., a temperature deviation from its set-point. However, in practice, limiting the power consumed by air conditioning devices has been shown to cause inadvertent load synchronization in the aggregate power consumption signal. Because thermal energy is stored in the home much like that of a leaky battery, during a demand response, this available energy is, in essence, consumed.

Throughout the home's limited power usage, temperatures will undoubtedly rise. Upon reinstatement of power, the governing controller will begin correcting this temperature deviation. Therefore, the HVAC system will run for a more extended period to restore temperature. However, on a macro-scale, because all HVAC systems have power reinstated in a relatively short time of one another, a significant spike in aggregate power consumption is experienced by the grid. Often this spike in aggregate power is worse than the events leading to the demand response event in the first place. Because peak-load is expensive and inefficient to supply, electric utilities employ demand response events to carefully distribute the affect of load synchronization. In this research, we are interested in the load synchronization effects caused by a demand

response event. Therefore, we employ a worst-case scenario for distributed SSHP and VSHP devices. The goal is to develop a control framework that mitigates load synchronization caused by a demand response event using sparse communication between neighboring homes. This task is accomplished via a properly structured MPC controller.

The demand response event employed in this study is defined as follows. During a simulated day, a population of residencies are subjected to an outside temperature disturbance, T_o . This outside temperature, T_o , follows a typical profile of a warm July summer's day. Throughout this outside temperature profile, the point in time, k , which coincides with the greatest outside temperature, $T_{o,max}$, is chosen to be the midpoint of the demand response event,

$$k^* = \arg \max_k T_o(k). \quad (2.24)$$

During this demand response period, Δt_{DR} , all compressors are shut down, i.e. $P_{agg}^k := \sum_{p=1}^N \eta_p m_p^k = 0 \forall k \in \{[k^* - \Delta t_{DR}/2], \dots, [k^* + \Delta t_{DR}/2]\}$. Although this demand response event is impractical in the sense of regulating peak power consumption and resembles more of a power outage, it does play an important role for observing the system response. Qualitative metrics are developed in Chapter 4 to compare the relative load synchronization effects of each controller including the lower-level controllers defined earlier. Moving forward, each case study will be simulated using the demand response event of Equation (2.24).

CHAPTER 3:

CONTROLLER

In this section, four optimal control frameworks are developed to efficiently schedule the control actions of a population of HVAC systems. Before these four control frameworks are defined, it helps to identify the design constraints and goals that influence each optimization program's structure. Firstly, and most importantly, these optimal control frameworks must regulate indoor air temperatures for the entire population of residential homes. With this cardinal feature, these optimization programs must balance the often competing aggregate power consumption constraints that benefit electric utilities. More precisely, these control frameworks must mitigate peak power consumption and, in general, reduce aggregate power consumption.

Moreover, under load synchronization inducing events like the demand response described in Section 2.6, these TCLs must distribute their control actions in such a way that mitigates extended oscillations. The combination of these design requirements serves as the foundation upon which our optimal control frameworks are built. As shown in Figure 3.1, four control frameworks are developed that are further classified by their control type and network topology. As it relates to the control type, a different MPC controller is dedicated to SSHPs and VSHPs. For each of these control types, a decentralized and centralized communication structure is used. For the decentralized communication structure, the network topology described in Section 2.2

		Network Topology	
		Centralized	Decentralized
Control Input	Continuous	$m \in [0, 1]$ $\mathbf{M} = \pi(\mathcal{N})$	$m \in [0, 1]$ $m_p = \pi(\mathcal{N}_p)$
	Discrete	$m \in \{0, 1\}$ $\mathbf{M} = \pi(\mathcal{N})$	$m \in \{0, 1\}$ $m_p = \pi(\mathcal{N}_p)$

Figure 3.1 Four control methods explored, where $\mathbf{M} = [m_1, \dots, m_N]^\top$.

is implemented. Otherwise, the centralized control framework's topological network structure relies on an omniscient entity like an electric utility to determine the control actions for its population. Both SSHP and VSHP centralized command structures are used as a comparative metric for the decentralized control framework. With the aid of network communication, the decentralized control frameworks will have comparable results as their centralized variants. The rest of this chapter is divided into four main sections: Decentralized Discrete Controller, Centralized Discrete Controller, Decentralized Continuous Controller, and Centralized Continuous Controller. For brevity,

Figure 3.2 System Diagram.

each of these control frameworks is further abbreviated as DD, CD, DC, and CC respectfully. The term “discrete” refers to the possible control actions of the SSHP i.e. $m_p^k \in \{0, 1\} \forall p, k \in \mathcal{N}, \{0, \dots, K\}$. Alternatively, the term “continuous” refers to the control action of VSHPs where $m_p^k \in [0, 1] \forall p, k \in \mathcal{N}, \{0, \dots, K\}$. An MPC solves for the control actions, $m^{i_k} \forall i_k \in \underline{H}$, where $\underline{H} := \{1, \dots, H\}$ represents a finite time horizon symbolically beginning at the k^{th} index.

Two indices are used to denote time, namely k and i_k . The first index, k , represents the simulation’s current time-step, whereas i_k is the current time-step of the simulated

horizon. It is helpful to parse these two time-steps because one deals with the past and present simulated values, while the other deals with future predictions.

The control action, m , is solved by minimizing an optimization program's objective function such that the model dynamics are satisfied. Once calculated, the first control action, $m^{i_k=1}$, is transmitted to the physical TCL device, also referred to as the plant. After the plant acts upon the control input, the response of the plant is then measured. As the state is observable, they are returned to the optimization program as an initial constrain for the subsequent iteration. The process then repeats itself by calculating the next time-step's control action, m^{k+1} . Inevitably, this cycle repeats until a termination condition is met. Features like the lockout period, discussed in Section 2.5, are built into the optimization program as a constraint. We observed that the MPC framework provided the most significant level of flexibility when constructing an optimal control algorithm for TCLs.

3.1 Decentralized Discrete Controller (DD)

Beginning with the decentralized discrete controller, each individual control action, m , is computed over a time horizon, $\underline{H}$, by way of a tuned optimization program. As depicted in Figure 3.2, the control block represents a Mixed Integer Quadratic

Program (MIQP), defined as,

$$\begin{aligned}
\min_{\mathbf{x}, m} \quad & J(k, \mathcal{X}_o) = \sum_{i_k \in \underline{\mathbb{H}}} \ell_{DD}(i_k, \mathbf{x}^{i_k}, m^{i_k}) \\
\text{s.t.} \quad & \mathbf{x}^1 = \mathcal{X}_o, \\
& \mathbf{x}^{i_k+1} = f_{DD}(i_k, \mathbf{x}^{i_k}, m^{i_k}, \theta^k), \\
& m^{i_k} \in \{0, 1\}, \\
& \forall i_k \in \underline{\mathbb{H}}.
\end{aligned} \tag{3.1}$$

In Equation (3.1), the objective penalty function, $\ell_{DD}(\cdot)$, is defined as,

$$\ell_{DD}(i_k, \mathbf{x}^{i_k}, m^{i_k}) = -\alpha \left(m^{i_k} - \frac{\sum_{j \in \mathcal{N}_p} m_j^{i_k}}{|\mathcal{N}_p|} \right)^2 - \beta \eta m^{i_k} + \gamma \tilde{T}^{i_k} + \zeta (\dot{T}_A^{i_k})^2, \tag{3.2}$$

and the dynamic model is represented as a set of linear constraints between incremented state values and the discrete state space representation of Equation (2.8), $\{\mathbf{x}^{i_k} \mid \mathbf{x}^1 = \mathcal{X}_o, \mathbf{x}^{i_k+1} = f_{dd}(i_k, \mathbf{x}^{i_k}, m^{i_k}, \theta^k) \forall i_k \in \underline{\mathbb{H}}\}$. The dynamic model, $f_{DD}(\cdot)$, is defined as,

$$f_{DD}(i_k, \mathbf{x}^{i_k}, m^{i_k}, \theta^k) = (I + \Delta t A) \mathbf{x}^{i_k} + \Delta t B \left(\eta m^{i_k} + \theta_4 \dot{T}_o^{i_k} + \theta_3 T_o^{i_k} \right). \tag{3.3}$$

From Equation (3.3), its observed that new state values are influenced by the control action, m^{i_k} , and the thermal parameters, θ^k , which are updated via the RLS algorithm defined in Section 2.4. To ensure continuity between simulated time-steps, initial state values, $\mathbf{x}^1$, are assigned as the measured response of the plant, $\mathcal{X}_o$. Lastly, a binary constraint is added to the all control decision variables so as not to violate the compressor's functionality.

The desired response of the control algorithm is determined by the cumulative objective terms within Equation (3.2). As is, $\ell_{DD}(i_k, \cdot)$ represents the objective penalty accumulated using it's respective time-steps' decision variables. The objective penalty, $\ell_{DD}(i_k, \cdot)$, is summed over the simulated horizon, $\underline{H}$ to form the objective program's total objective value, $J(\cdot)$. The goal of any optimization program is the minimization or maximization of an objective function such that a desired response is obtained. Our desired response is the proper control of temperature and aggregate power consumption. The desired response thereby dictates the structure of the optimization program. With computational efficiency in mind, the objective program of Equation (3.2) is composed of linear and quadratic terms. Each term is linked to a particular attribute of the desired response. The first objective penalty term,

$$-\alpha \left(m^{i_k} - \frac{\sum_{j \in \mathcal{N}_p} m_j^{i_k}}{|\mathcal{N}_p|} \right)^2, \quad (3.4)$$

aims to reduce load synchronization effects by maximizing the difference in control action of the p^{th} SSHP and average control action of its neighbor set, $\mathcal{N}_p$. Intuitively when the average control action of a SSHP's neighbor set registers high, the optimizer will likely turn $m_p = 0$. Alternatively, when the average control action of a neighbor-set registers low, the optimizer will likely turn $m_p = 1$. Because of this, a damping effect is experienced in the aggregate power consumption signal. In addition to reducing load synchronization, this objective term also minimizes peak power consumption by virtue of its smoothening effect.

The second term objective penalty, $-\beta \eta m^{i_k}$, also influences aggregate power consumption. Unlike the first term, this objective penalty helps minimize individual

power consumption, resulting in a reduction of aggregate power consumption. Because no singular agent has omniscient knowledge of its population-sets' control actions, $\mathcal{N}$, the resulting aggregate power consumption is deemed sub-optimal. However, as shown in Chapter 4, it does provide comparable results to the following centralized control framework of Section 3.2. For clarity's sake, heat is removed from the residential control volume via η . By convention $\eta \leq 0$, therefore to minimize power the product of $-\beta\eta m^{i_k} \geq 0$ must be positive.

The third objective penalty term, $\gamma\tilde{T}^{i_k}$, is responsible for maintaining indoor air temperatures at or near its set-point temperature, T_{sp} . This objective penalty maintains indoor air temperature via a soft constraint. In doing so, temperature allowed to deviate from the dead-band, thereby reducing short cycling. In this proposed objective term, a penalty is linearly accrued when the temperature deviates above or below its respective dead-band threshold, that being δ^+ and δ^- . The term $\tilde{T}^{i_k}$ is a double-hinged function which is mathematically expressed as,

$$\tilde{T}^{i_k} = \max(\delta^- - T_A^{i_k}, 0, T_A^{i_k} - \delta^+).$$

With this formulation, the indoor air temperature, $T_A^{i_k}$, is observed to not incur a cost when between δ^+ and δ^- . This allows the optimizer to minimize power related objective terms all the while still being commanded between the dead-band threshold. This double hinge function is graphically depicted in Figure 3.3.

The last objective penalty term, $\zeta(\dot{T}_A^{i_k})^2$, focuses on reducing the rate at which $T_A^{i_k}$ changes between time-steps. This penalty term, in the context of ON/OFF control is deemed less important although still necessary in the event of rapid temperature changes.

Figure 3.3 Objective soft constraint $\tilde{T}$ calculated using double hinge function.

Each objective penalty term is accompanied by an objective constant, those being α , β , γ , and ζ . These objective constants are responsible for tuning the desired response of the optimization program. Earlier in Section 2.6 we alluded that an electric utilities' power requirements need to be balanced with the end-users comfort. If indoor temperatures are maintained closer to their desired reference temperature, assuming short cycling is avoided, the grid will inevitably experience more power consumption. Concessions ought to be made via the indoor temperature to reduce aggregate power consumption. Thermodynamics aside, these objective constants let the operator tune the controller's desired response to reach a pleasant equilibrium for the end-user.

When an objective constant is increased, it linearly scales its objective penalty term, thereby increasing the cost accrued for the same decision variables. If all

objective constants are raised similarly, the controller's response will experience no change. Therefore, change in the desired response comes about from the relative magnitudes of each objective constant term. For instance, if one were to increase β a reduction in aggregate power consumption, P_{agg} , will occur. It should be noted these constants are positively valued real numbers.

The entirety of this section mathematically represents the control and model blocks of Figure 3.2. Once the optimizer calculates a horizon of control inputs, $m^{i_k} \forall i_k \in \underline{H}$, the control architecture then sends the first control input to the plant. Then the state variables are assigned as $\mathcal{X}_o$ for the following iteration. at each time-step, k , thermal parameters which characterize the ETP model are then updated via Algorithm (1).

3.2 Centralized Discrete Controller (CD)

The next control framework, similar to the DD controller described in Section 3.1, uses a MIQP to compute the control actions for an entire population of SSHP. The main difference is how control actions are calculated. For the DD control framework, each home is responsible for computing its control action. Moreover, each TCL uses sparse communication to compute the average control action of the p^{th} agent's neighbor-set $\mathcal{N}_p$.

With the CD variant, a single entity computes the entire population's control action. Necessarily, this feature requires that the central entity, possibly the electric utility, has access to all state and control inputs of its population set. Due to its clear omniscience, the governing control framework is deemed to compute better solution than its decentralized counterpart because it possesses and thereby acts upon the

entire population set's information. The CD controller's MIQP is defined as,

$$\begin{aligned}
\min_{\mathbf{X}, \mathbf{M}} \quad & J(k, \mathcal{X}_o) = \sum_{i_k \in \underline{\mathbf{H}}} g_{CD}(i_k, \mathbf{M}^{i_k}) + \ell_{CD}(i_k, \mathbf{X}^{i_k}, \mathbf{M}^{i_k}) \\
\text{s.t.} \quad & \mathbf{x}_p^1 = \mathcal{X}_{o,p}, \\
& \mathbf{x}_p^{i_k+1} = f_{CD}(p, i_k, \mathbf{x}_p^{i_k}, m_p^{i_k}, \theta_p^k), \\
& m_p^{i_k} \in \{0, 1\}, \\
& \forall p \in \mathcal{N}, \forall i_k \in \underline{\mathbf{H}},
\end{aligned} \tag{3.5}$$

where $\mathbf{X}^{i_k} = [\mathbf{x}_1^{i_k \top}, \dots, \mathbf{x}_N^{i_k \top}]^\top$ and $\mathbf{M}^{i_k} = [m_1^{i_k}, \dots, m_N^{i_k}]^\top$ are the state and control variables concatenated in a convenient array. After Equation (3.5) computes $\mathbf{M}^{i_k}$, the centralized entity then distributes each control action to its respective SSHP. As shown in Equation (3.5), the cumulative objective function, $J(\cdot)$, is now conveniently separated into two functions, where $g_{CD}(\cdot)$ is defined as,

$$g_{CD}(i_k, \mathbf{M}^{i_k}) = \alpha \sum_{z=0}^{z_m} \left(\xi^{(z)} \left| \sum_{p \in \mathcal{N}} (m_p^{i_k-z} - m_p^{i_k-z-1}) \right| \right), \tag{3.6}$$

and $\ell_{CD}(\cdot)$ is defined as,

$$\ell_{CD}(i_k, \mathbf{X}^{i_k}, \mathbf{M}^{i_k}) = \sum_{p \in \mathcal{N}} -\beta \eta m_p^{i_k} + \gamma \tilde{T}_p^{i_k} + \zeta (\dot{T}_{A,p}^{i_k})^2. \tag{3.7}$$

Within Equation (3.5) the model dynamics are simulated as a set of equality constraints between incremented state values and the function, $f_{CD}(\cdot)$, where $f_{CD}(\cdot)$ is defined similarly as $f_{DD}(\cdot)$. However, now $f_{CD}(p, \cdot)$ is an explicit function of the p^{th} home. The number of equality constraints is linearly proportional to the number

of homes, N , being simulated.

$$f_{CD}(p, i_k, \mathbf{x}_p^{i_k}, m_p^{i_k}, \theta_p^k) = (I + \Delta t A) \mathbf{x}_p^{i_k} + \Delta t B \left(\eta m_p^{i_k} + \theta_4 \dot{T}_o^{i_k} + \theta_3 T_o^{i_k} \right). \quad (3.8)$$

The outside temperature, $T_o^{i_k}$, and its rate of change, $\dot{T}_o^{i_k}$, are assumed to be uniform, meaning each home of our simulated population experiences the same outside temperature. Stated differently, it is reasonable to assume these disturbance values are the same amongst agents. In a physical setting, this disturbance value is easily measured on-site and communicated to the governing controller.

By means of Equation 3.6, the newly formulated objective function takes full advantage of the population's control action, $\mathbf{M}^{i_k}$. The new communication term, $g_{CD}(\cdot)$, minimizes the difference in aggregate control effort between congruent time-steps. The aggregate control effort is defined as $m_{agg}^{i_k} = \sum_{p \in \mathcal{N}} m_p^{i_k}$. By finding the absolute difference in the aggregate control effort along a decremented sliding window, the optimizer can minimize deviations in peak aggregate power consumption. Because Equation (3.6) and (3.2) deal with the same attribute, that being minimizing load volatility, they are both multiplied by the same objective constant, α . The term ξ^z is used to exponentially decay each subsequent passing window. All other objective terms contained within $\ell_{CD}(\cdot)$, are similar to Equation (3.2), therefore their redefinition are omitted. The CD controller defined by Equation (3.5) was developed as a comparative tool. Its purpose is the validation of the DD controller defined by Equation (3.1). As it stands, this CD controller is unfit to regulate large populations of SSHPs. It does, however, provide both quantitative and qualitative measures to help evaluate the DD performance.

Unlike the MIQP of Equation (3.1), the number of decision variables scales lin-

early with the size of the population, N . From this, we deduce that the computational complexity rises combinatorially. There exists a population size, N , which will inevitably cause the optimizer to stall. As shown in Equation (3.5), this optimization program solves for the control actions, $\mathbf{M}$, over a time horizon, $\underline{H}$, for a population size, N . Not only does the number of decision variables increase, but also the number of model constraints and initial conditions. The difficulty of this CD controller is that all control inputs are discrete elements of the binary set, $m \in \{0, 1\}$. While solving, the optimizer must heuristically check a multitude of possible solutions. For this reason, the CD controller is unfit to compute the control actions for a large population size of TCLs.

Due to the discrete nature of the control input, $m_p \in \{0, 1\} \forall p \in \mathcal{N}$, the optimizer must employ a highly specialized solver to minimize its objective. The optimizer uses a branch and bound technique to relax integer constraints, approximate a solution, then repeat until a local/global minima is reached. This iterative method, unlike typical Quadratic Program (QP) is more costly to solve. For this reason, the simulation results generated by the CD controller will be considered a gold standard but not feasible in a natural setting. This concluded the control frameworks for the SSHP.

3.3 Decentralized Continuous Controller (DC)

This subsection discusses a decentralized control framework for VSHPs. Two key differences exist between the DD controller and this DC controller. As the name suggests the control action is continuous between being OFF and ON, $m \in [0, 1]$. This change allows the following optimization program to be defined as a typical QP. Secondly, as defined in Section 2.1.2, the VSHP's model use $\sigma := \dot{m}$ as the control input. This allows the optimizer to minimize the rate which the control action, m ,

changes. The optimization program graphically represented as the controller block in Figure 3.2, is defined as,

$$\begin{aligned}
\min_{\mathbf{x}, m, \sigma} \quad & J(k, \mathcal{X}_o) = \sum_{i_k \in \underline{\mathbb{H}}} \ell_{DC} (i_k, \mathbf{x}^{i_k}, m^{i_k}, \sigma^{i_k}) \\
\text{s.t.} \quad & \mathbf{x}^1 = \mathcal{X}_o, \\
& m^{i_k} \in [0, 1], \\
& m^{i_k+1} = m^{i_k} + \Delta t \sigma^{i_k}, \\
& \mathbf{x}^{i_k+1} = f_{DC} (i_k, \mathbf{x}^{i_k}, m^{i_k}, \sigma^{i_k}, \theta^k), \\
& \forall i_k \in \underline{\mathbb{H}}.
\end{aligned} \tag{3.9}$$

In Equation (3.9), an optimizer minimizes the objective function, $J(\cdot)$, with respect to the state and input decision variables, i.e., $\mathbf{x}$, m , and σ . By manipulating these particular decision variables, a horizon of control actions are calculated such that the model dynamics, initial conditions, and saturation limits are maintained. An equality constraint is placed between the initial state values, $\mathbf{x}^{i_k=1}$, and the measured/predicted response of the plant, $\mathcal{X}_o$ to maintain continuity between time-steps. Likewise, m is constrained between 0 and 1, thereby preserving its associated VSHP devices within safe operating conditions. Lastly, the model dynamics of Equation (2.10), is expressed as a set of linear equality constraints, where the function $f_{DC}(\cdot)$ is defined as,

$$f_{DC} (i_k, \mathbf{x}^{i_k}, m^{i_k}, \sigma^{i_k}, \theta^k) = (I + \Delta t A) \mathbf{x}^{i_k} + \Delta t B \left(\eta (m^{i_k} + \theta_5 \sigma^{i_k}) + \theta_4 \dot{T}_o^{i_k} + \theta_3 T_o^{i_k} \right). \tag{3.10}$$

As shown in Equation (3.1), the objective function, $J(\cdot)$, is the summation of each

time-step's cumulative objective penalty, $\ell_{DC}(\cdot)$, defined as,

$$\ell_{DC}(i_k, \mathbf{x}^{i_k}, m^{i_k}, \sigma^{i_k}) = \alpha \left(m^{i_k} - \frac{\sum_{j \in \mathcal{N}_p} m_j^{i_k}}{|\mathcal{N}_p|} \right)^2 - \beta \eta m^{i_k} + \gamma \tilde{T}^{i_k} + \zeta (\dot{T}_A^{i_k})^2 + \tau (\sigma^{i_k})^2. \quad (3.11)$$

Again, each objective penalty term of Equation (3.11) is dedicated to a certain attribute of the desired response. Based the preferences of electric utility service providers and end-users alike, this DC framework must maintain indoor air temperatures while simultaneously minimizing the aggregate power consumption signal and its associated load volatility.

The first objective penalty term,

$$\alpha \left(m^{i_k} - \frac{\sum_{j \in \mathcal{N}_p} m_j^{i_k}}{|\mathcal{N}_p|} \right)^2, \quad (3.12)$$

is similar to Equation (3.4). However, now the goal is to minimize the difference between the p^{th} agent's control action and the average control action of its neighbor set $\mathcal{N}_p$. Intuitively, the optimizer will try to group control actions. For example, in the middle of a warm July day, when a significant portion of VSHPs are cooling at some percentage of max capacity, this objective term will tend to cause neighboring homes to output similar control actions. This feature prevents rapid or sudden changes to the control effort of any one system. Moreover, depending on its size, when the demand response event restores power, the average control effort will effectively drop. This causes a gradual increase back to normalacy. Lastly, this term reduces load synchronization's effects caused by internal/external measures. For context, load synchronization is defined as the grouping of control actions such that oscillatory

ringing is experienced in the aggregate power consumption signal.

The objective penalty terms $-\beta\eta m^{i_k}$, $\gamma\tilde{T}^{i_k}$, and $\zeta(\dot{T}_A^{i_k})^2$ of Equation 3.11, faithfully match that of the DD and CD control frameworks. Here, $-\beta\eta m^{i_k}$ minimizes individual power consumption, $\gamma\tilde{T}^{i_k}$ maintains indoor air temperatures, T_A within their respective dead-band threshold, and $\zeta(\dot{T}_A^{i_k})^2$ reduces the rate at which indoor air temperature can change. In the case of the former DD and CD MPC controllers, minimizing a home's temperature rate of change for discrete control actions had little effect on the response of the plant. However, in the DC framework, now not only is temperature maintained at its set-point but also prevented from rapid change rates.

The last objective term, $\tau(\sigma^{i_k})^2$, minimizes the rate at which a VSHP cools the interior space of a home. For instance, after a demand response event, once power is reinstated, the controller's logical action is to minimize temperature deviations above the upper dead-band. Necessarily this means that VSHP increase its cooling capacity. Inadvertently, this leads to gradual changes in the aggregate power consumption signals. In essence, this term helps smooth any and all variations in aggregate power consumption.

3.4 Centralized Continuous Controller (CC)

Finally, a CC control framework is developed for comparison purposes against its DC VSHP variant. This control framework aims to identify the controller response, given omniscient knowledge of the state and control variables of its population. Acting as a benchmark, the simulation results generated by this control framework gives a sense of how well the DC controller compares. Instead of communicating nodes, this CC is structured such that all independent agents transmit previous state and control actions to the entity responsible for calculating new control actions, i.e., an

electric utility. Once calculated, the utility then sends new control actions to their respective VSHP where it is then acted upon by the plant. After the plant's state response is measured/predicted, the cycle then repeats. The QP defining this CC control framework is further defined as,

$$\begin{aligned}
\min_{\mathbf{X}, \mathbf{M}, \sigma} \quad & J(k, \mathcal{X}_o) = \sum_{i_k \in \underline{\mathbf{H}}} \ell_{CC} (i_k, \mathbf{X}^{i_k}, \mathbf{M}^{i_k}, \sigma^{i_k}) \\
\text{s.t.} \quad & \mathbf{x}_p^1 = \mathcal{X}_{o,p}, \\
& m_p^{i_k} \in [0, 1], \\
& m_p^{i_k+1} = m_p^{i_k} + \Delta t \sigma_p^{i_k}, \\
& \mathbf{x}_p^{i_k+1} = f_{CC} (p, i_k, \mathbf{x}_p^{i_k}, m_p^{i_k}, \sigma_p^{i_k}, \theta_p^k), \\
& \forall p \in \mathcal{N}, \forall i_k \in \underline{\mathbf{H}},
\end{aligned} \tag{3.13}$$

where $\mathbf{X}^{i_k} = [\mathbf{x}_1^{i_k}, \dots, \mathbf{x}_N^{i_k}]^\top$ and $\mathbf{M}^{i_k} = [m_1^{i_k}, \dots, m_N^{i_k}]^\top$ are defined as the collection of state and control decision variables at the i_k^{th} time-step. As may become apparent, the number of decision variables is directly proportional to the size of the population being simulated. Due to this increased number of decision variables, the associated complexity of the controller also rises. Both the CD and CC controllers, described by Equations (3.5) and (3.13) respectfully, help compare the simulation results of their decentralized counterparts.

In Equation (3.13), the dynamic model is represented as a set of linear constraints between incremented state values and the model dynamic function $f_{CC}(\cdot)$, initialized

with the state response of the plant, $\mathcal{X}_o$. The function, $f_{CC}(\cdot)$, is further defined as,

$$\begin{aligned} f_{CC}(p, i_k, \mathbf{x}_p^{i_k}, m_p^{i_k}, \sigma_p^{i_k}, \theta_p^k) &= (I + \Delta t A) \mathbf{x}_p^{i_k} \\ &+ \Delta t B \left(\eta (m_p^{i_k} + \theta_5 \sigma_p^{i_k}) + \theta_4 \dot{T}_o^{i_k} + \theta_3 T_o^{i_k} \right). \end{aligned} \quad (3.14)$$

Next, the objective function, $\ell_{CC}(\cdot)$, has the form,

$$\begin{aligned} \ell_{CC}(i_k, \mathbf{X}^{i_k}, \mathbf{M}^{i_k}, \sigma^{i_k}) &= \alpha \left| \sum_{p \in \mathcal{N}} (m_p^{i_k+1} + m_p^{i_k}) \right| \\ &+ \sum_{p \in \mathcal{N}} \left(-\beta \eta_p m_p^{i_k} + \gamma \tilde{T}_p^{i_k} + \zeta (\dot{T}_{A,p}^{i_k})^2 + \tau (\sigma_p^{i_k})^2 \right). \end{aligned} \quad (3.15)$$

In Equation (3.15), The first objective term,

$$\alpha \left| \sum_{p \in \mathcal{N}} (m_p^{i_k+1} + m_p^{i_k}) \right|, \quad (3.16)$$

minimizes the absolute difference in aggregate control effort between each simulated horizon time-step, i_k . As to be expected, the last four objective penalty terms $-\beta \eta_p m_p^{i_k}$, $\gamma \tilde{T}_p^{i_k}$, $\zeta (\dot{T}_{A,p}^{i_k})^2$, and $\tau (\sigma_p^{i_k})^2$ are similar to the DC control framework. However, the cumulative objective penalty, $J(p, \cdot)$, is now summed over the population-set, $\mathcal{N}$. For brevity, their redefinitions will be omitted. Again, each objective term defined in Equation (3.15) is multiplied with an objective constant, that being α , β , γ , ζ , and τ . These objective constants give credence to the optimizer to selectively control particular attributes of the objective function.

With complete state and control knowledge, our goal is to determine what a more optimal solution might look like and use it as a comparative metric against the DC controller.

CHAPTER 4:

CASE STUDY

In Chapter 3, four optimal control frameworks were developed to systematically schedule the control actions of a population TCL devices. These optimal control frameworks help maintain indoor air temperatures while having beneficial aggregate power characteristics for both SSHPs and VSHPs alike. Of these two TCL devices, a centralized and decentralized controller have been proposed.

Using the MPC algorithm graphically depicted in Figure 3.2, each of these four control frameworks are simulated. At the machine level, the three controllers discussed in Chapter 2 are also used. These lower-level controllers are the hysteresis, advanced hysteresis, and PID controllers defined by Equations (2.5), (2.12), and (2.15), respectfully. Both the hysteresis and advanced hysteresis controllers are designed to regulate SSHP devices, whereas, the PID controller regulates VSHP devices.

As previously stated, T_o and its derivative $\dot{T}_o$ represent disturbances to the thermal system. Because the control frameworks proposed in Chapter 3 use a receding horizon, outside temperature data must be known over the entire simulated horizon $\underline{H}$. In our case, this horizon lasts on the order of tens of minutes. Given modern meteorological forecasting techniques, it is reasonable to assume this future outside temperature data is known at the time of simulation with sufficient accuracy. The outside temperatures rate of change, $\dot{T}_o$, is then determined using numerical differentiation.

In reality, two likely scenarios exist. Outside temperature data is adaptively sampled at the simulator’s frequency or is sampled at fixed intervals where internal points are then interpolated. In a simulation, like the ones presented in this section, the latter is chosen. Outside temperature data is acquired from the National Solar Radiation Database (NSRDB) in the form of a Typical meteorological Year (TMY) (Sengupta et al., 2018). A TMY dataset contains, amongst other qualitative properties, ambient outside temperature data sampled at an hourly rate. As the name suggests, this dataset contains the most usual weather conditions for a given region and is well suited for our application of weather prediction.

HVAC units are typically sold in half-ton increments, where one ton of cooling is defined as the amount of heat required to freeze/melt two-thousand pounds of water in a twenty-four-hour period ($12,000Btu/hr$). Before installing an HVAC unit, the discrete tonnage is chosen according to the thermodynamic properties of the space it is required to condition.

In this study, an air conditioning system’s cooling capacity, $\eta < 0$, is determined based on the time needed to cool that home’s indoor air temperature, T_A , from its upper dead-band, δ^+ , to its lower dead-band, δ^- . The term $\eta(\cdot)$ is a function of a residential home’s thermodynamic properties. In our case, the thermodynamic properties are sufficiently characterized by the parameter vector, θ .

We created a population of heterogeneous residential homes. This is done so by generating thermal parameter values about a known statistical distribution. Our statistical distribution is Gaussian and is fully characterized by the mean and standard deviation values listed in Table 4.1. Once a population-set’s thermal characteristics are determined, HVAC tonnage is then calculated. A cycle duration of thirty minutes

Table 4.1 Thermal Parameters.

Thermal Parameter	Mean	Std.
$C_A \left[\frac{Btu}{^\circ F} \right]$	1,080	54
$C_M \left[\frac{Btu}{^\circ F} \right]$	4,280	214
$\frac{1}{R_1} \left[\frac{Btu}{^\circ F \cdot hr} \right]$	520	26
$\frac{1}{R_2} \left[\frac{Btu}{^\circ F \cdot hr} \right]$	7,050	353

Table 4.2 Simulation parameters

	(H)	(AH)	(PID)	(DD)	(CD)	(DC)	(CC)
	4.1.1	4.1.2	4.1.3	4.4	4.6	4.8	4.10
α	-	-	-	300	1,000	200	300
β	-	-	-	0.001	0.001	0.001	0.001
γ	-	-	-	5,000	1,200	5,000	8,000
ζ	-	-	-	0.01	-	100	100
τ	-	-	-	-	-	1,000	1,500
Homes (N)	1,000			50			
Time-step (K)	4,000						
Horizon (H)	-	-	-	20			
Step-length (Δt)	23.4 [sec]						
D.R. length	20 [min]						

at $95^\circ F$ was chosen when computing η . Besides thermal parameters listed in Table 4.1, all other simulation parameters are listed in Table 4.2.

4.1 Baseline Simulation Results

Each of the following baseline studies are accompanied by a sudo algorithm defining how its corresponding simulation results are generated. Each algorithm has the following sequence of events; first parameters are generated, then both the time-step, k , and agent, p are iterated. During each iteration, the governing control algorithm determines the control action, m . This value is sent to the plant where it is then acted

upon. The simulated response and control actions are then stored. Finally after the final iterate is completed, graphical plots are then generated.

4.1.1 Hysteresis Controller

Algorithm 2 Hysteresis Control Sequence

```

Initialize  $\mathcal{P} \quad \emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$    Generate  $\theta_0$ 

 $\mathcal{D}$    Set Initial Conditions

for  $k = 1$  to  $K$  do
    for  $p = 1$  to  $N$  do
         $m_p^k \quad f(\mathcal{D}, k, p)$  Equation (2.5)
         $x_p^{k+1} \leftarrow Plant(\mathcal{P}, \mathcal{D}, k, p, m_p^k)$  Equation (2.8)
         $\mathcal{D} \quad x_p^{k+1}, m_p^k$ 
    end for
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.1 Hysteresis controller results generated via Algorithm (2)($N=1,000$).

Using the hysteresis toggle condition of Equation (2.5), Algorithm (2) generates the simulation results presented in Figure 4.1. Load aggregation is observed in the aggregate power consumption subplot of Figure 4.1b. This claim is further justified by the secondary frequency peak of the Fast Fourier Transform (FFT) subplot of Figure 4.1c. During both the morning hours and after the demand response event, indoor temperatures are observed to group up in distinct bands before being slowly diversified. This, and the following PID control framework of Section 4.1.3 represent the only uncoordinated control architectures. uncoordinated

4.1.2 Advanced Hysteresis Controller

Algorithm 3 Advanced Hysteresis Control Sequence

```

Initialize  $\mathcal{P}$      $\emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$     Generate  $\theta_0$ 

 $\mathcal{D}$     Set Initial Conditions

for  $k = 1$  to  $K$  do
     $\tilde{m}^k$      $g(\mathcal{D})$  Equation (2.11)

    for  $p = 1$  to  $N$  do
         $m_p^k$      $f(\mathcal{D}, k, p, \tilde{m}_p^k)$  Equation (2.12)
         $x_p^{k+1} \leftarrow Plant(\mathcal{P}, \mathcal{D}, k, p, m_p^k)$  Equation (2.8)
         $\mathcal{D}$      $x_p^{k+1}, m_p^k$ 
    end for
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.2 Advanced hysteresis controller results generated via Algorithm (3) ($N=1,000$).

By using the $\tilde{m}$ controller of Equation (2.12) within Algorithm (3), simulation results presented in Figure 4.2 are generated. Minor load aggregation is still observed during both the morning and after the demand response event. However, the load aggregation's amplitude and rate of decay is significantly improved. Temperatures are evenly distributed, but temperatures tend to deviate further above the upper dead-band when compared to the hysteresis controller above.

4.1.3 PID Controller

Algorithm 4 PID Control Sequence

```

Initialize  $\mathcal{P} \quad \emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$    Generate  $\theta_0$ 

 $\mathcal{D}$    Set Initial Conditions

for  $k = 1$  to  $K$  do
    for  $p = 1$  to  $N$  do
         $e_p^k \quad h(\mathcal{D}, k, p)$  Equation (2.14)
         $u_p^k \leftarrow g(e_p^k, k, p)$  Equation (2.15)
         $m_p^k \leftarrow f(u_p^k)$  Equation (2.16)
         $x_p^{k+1} \quad Plant(\mathcal{P}, \mathcal{D}, k, p, m_p^k)$  Equation (2.10)
         $\mathcal{D} \quad x_p^{k+1}, m_p^k, e_p^k$ 
    end for
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.3 Advanced hysteresis controller results generated via Algorithm (4) ($N=1,000$).

In Algorithm (4), Equations (2.14) and (2.15) are simulated using a zero-order-hold to discretize newly computed control actions. Of the three baseline studies, this PID framework has the smoothest aggregate power consumption subplot as observed in Figure 4.3b.

4.2 Optimal Control Simulation Results

In this section, simulation results for the four optimal control frameworks discussed in Chapter 3 are observed. Each control framework is accompanied by an algorithm which generates integrated state and estimated thermal parameter values. These four control frameworks all require an optimization program to be solved. All simulations are performed with the Julia programming language. Moreover, within Julia, all optimization programs are solved via JuMP (Dunning et al., 2017), Julia’s open source optimization package.

By way of Algorithm (1), elements of the thermal parameter vector, θ , are systematically updated to improve both the accuracy and resilience of the predictive model. These incremental updates provide the ETP model the ability to accurately mimic the dynamics of its plant. In the following simulations, an error signal is generated for each element of the thermal parameter vector. This error signal is defined as the difference between the plant and model parameter values, i.e. $\theta^p - \theta^m$. Where plant parameters are statistically generated about the mean and standard deviation values listed in Table 4.1. Once generated, model parameters are then generated about their corresponding plant parameter with noise included to simulate initial measurement error.

4.2.1 Decentralized Discrete Controller

Algorithm 5 Decentralized Discrete MPC Sequence

```

Initialize  $\mathcal{P} \quad \emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P} \quad$  Generate  $\theta_0$ 

 $\mathcal{D} \quad$  Set Initial Conditions

for  $k = 1$  to  $K$  do
    for  $p = 1$  to  $N$  do
         $m_p^{i_k \in \underline{H}} \quad Controller(\mathcal{D}, \mathcal{P}, k, p)$  Equation (3.1)
         $x_p^k \leftarrow Plant(\mathcal{D}, \mathcal{P}, k, p, m_p^k)$  Equation (2.8)
         $\mathcal{D} \quad x_p^k, m_p^{i_k \in \underline{H}}$ 
    end for
    if  $k \geq 2$  then
         $\mathcal{D} \quad Lockout(\mathcal{D})$ 
         $\mathcal{P} \quad$  Algorithm (1)
    end if
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.4 DD controller results generated via Algorithm (5)

(N=50).

Figure 4.5 DD estimated thermal parameter values represented as error signal (N=50).

Using the optimization program of Equation (3.1) within Algorithm (5) simulation data is generated then presented Figures 4.4 and 4.5.

4.2.2 Centralized Discrete Controller

Algorithm 6 Centralized Discrete MPC Sequence

```

Initialize  $\mathcal{P} \leftarrow \emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$     Generate  $\theta_0$ 

 $\mathcal{D}$     Set Initial Conditions

for  $k = 1$  to  $K$  do
     $\mathbf{M}^{i_k \in \underline{\mathbb{H}}} \leftarrow \text{Controller}(\mathcal{D}, \mathcal{P}, k, p)$  Equation (3.5)
     $\mathbf{X}^k$      $\text{Plant}(\mathcal{D}, \mathcal{P}, k, p, \mathbf{M}^k)$  Equation (2.8)
     $\mathcal{D} \leftarrow \mathbf{X}^k, \mathbf{M}^{i_k \in \underline{\mathbb{H}}}$ 
    if  $k \geq 2$  then
         $\mathcal{D} \leftarrow \text{Lockout}(\mathcal{D})$ 
         $\mathcal{P} \leftarrow \text{Algorithm (1)}$ 
    end if
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.6: CD controller results generated via Algorithm (6)

(N=50).

Figure 4.7 CD estimated thermal parameter values represented as error signal ($N=50$).

The CD simulation results shown in Figures 4.6 and 4.7 are generated using Algorithm (6). Distinct level sets are observed in the aggregate power consumption subplot of Figure 4.6b. This is likely caused by objective programs reluctance to change aggregate control efforts between time-steps.

4.2.3 Decentralized Continuous Controller

Algorithm 7 Decentralized Continuous MPC Sequence

```

Initialize  $\mathcal{P} \leftarrow \emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$     Generate  $\theta_0$ 

 $\mathcal{D}$     Set Initial Conditions

for  $k = 1$  to  $K$  do
    for  $p = 1$  to  $N$  do
         $m_p^{i_k \in \underline{\mathbf{H}}}$      $Controller(\mathcal{D}, \mathcal{P}, k, p)$  Equation (3.9)
         $x_p^k \leftarrow Plant(\mathcal{D}, \mathcal{P}, k, p, m_p^k)$  Equation (2.10)
         $\mathcal{D} \leftarrow x_p^k, m_p^{i_k \in \underline{\mathbf{H}}}$ 
    end for
    if  $k \geq 2$  then
         $\mathcal{P} \leftarrow Algorithm(1)$ 
    end if
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.8 DC controller results generated via Algorithm (7)

(N=50).

Figure 4.9 DC estimated thermal parameter values represented as error signal (N=50).

Like the PID controller, continuous control actions are computed via the QP of Equation 3.9. Algorithm (7) is used to generate the simulation results shown in Figures 4.8 and 4.9. Of the baseline and discrete controllers discussed, the DC and CC frameworks have the smoothest aggregate power consumption. A key advantage of this approach is the smooth aggregate power consumption ramp rate after the simulated demand response event. Likewise, a singular FFT frequency peak is observed with minimal high frequency noise. The estimated thermal parameters of Figure 4.9 shows minor difficulty when converging to their true parameters values.

4.2.4 Centralized Continuous Controller

Algorithm 8 Centralized Continuous MPC Sequence

```

Initialize  $\mathcal{P}$      $\emptyset, \mathcal{D} \leftarrow \emptyset$ 

 $\mathcal{P}$     Generate  $\theta_0$ 

 $\mathcal{D}$     Set Initial Conditions

for  $k = 1$  to  $K$  do
     $\mathbf{M}^{i_k \in \underline{\mathbb{H}}} \leftarrow \text{Controller}(\mathcal{D}, \mathcal{P}, k, p)$  Equation (3.13)
     $\mathbf{X}^k$      $\text{Plant}(\mathcal{D}, \mathcal{P}, k, p, \mathbf{M}^k)$  Equation (2.10)
     $\mathcal{D}$      $\mathbf{X}^k, \mathbf{M}^{i_k \in \underline{\mathbb{H}}}$ 
    if  $k \geq 2$  then
         $\mathcal{P}$     Algorithm (1)
    end if
end for

```

(a) Indoor air temperatures.

(b) Aggregate power consumption signal.

(c) Fast Fourier transform of aggregate power consumption signal.

Figure 4.10 CC controller results generated via Algorithm (8) ($N=50$).

Figure 4.11 CC estimated thermal parameter values represented as error signal (N=50).

Like the CD framework, a population of control actions are simultaneously calculated via the QP of Equation (3.13). Algorithm (8) is used to generate the simulation results shown in Figures 4.10 and 4.11. The CC and DC control frameworks are observed to have similar thermal parameter convergence properties.

4.3 Comparison

Several quantitative metrics of performance are now introduced. These metrics help compare the simulation results generated via each control framework and help serve as a point of reference for future studies. They help relate the temperature and aggregate power demand requirements discussed in Chapter 2.

A first metric is introduced to measure the cumulative time a population spends above its upper dead-band threshold, δ^+ . Next, three additional metrics are introduced to relate aggregate power performance. These power metrics, denoted P_i , P_f , and E_T , provide a relative measurement to compare varying control frameworks. The terms P_i and P_f are both ratios between the total consumable power and the maximum power consumed before and after our demand response event. While P_i and P_f are measures of relative power, E_T is a ratio the systems exergy and the total aggregate energy expended during simulation. These terms are expressed as ratios to more intrinsically represent each control type's response, thereby allowing a comparison between baseline and optimal control frameworks. Lastly, to gain insight into the phenomenon of load aggregation, we apply an FFT on the aggregate power consumption signal. An FFT decomposes a time-domain signal into its corresponding frequency domain. A one-to-one representation of the amplitude of the original signal at each frequency is then plotted. If the FFT plot has significant peaks at a frequency away from zero, we conclude that repeated load synchronization at this frequency can

be observed in the original signal.

Table 4.3 Quantitative Performance Metrics.

Section	Controller Type	t_{dev} [hrs/TCL]	P_i	P_f	E_T
4.1.1	Hysteresis	0.81	79.7%	73.6%	15.2%
4.1.2	Advanced Hysteresis	3.05	42.5%	49.0%	15.0%
4.1.3	PID	0.80	35.9%	61.4%	14.5%
4.2.1	Decentralized Discrete	0.36	44.1%	50.0%	16.6%
4.2.2	Centralized Discrete	3.04	38.4%	56.5%	16.3%
4.2.3	Decentralized Continuous	0.72	36.1%	30.8%	14.4%
4.2.4	Centralized Continuous	1.02	35.1%	31.1%	14.4%

Based on the simulation results of Sections 4.1 - 4.2 and the performance metrics listed in Table 4.3 several observations are made. Firstly, in Figure 4.1b the phenomenon of load synchronization is observed. This claim is further exemplified by the prominent peaks of Figure 4.1c. As observed in the aggregate power consumption subplot, a ringing occurs during the simulated day's morning hours. This ringing continues until being damped out by the heterogeneous thermal characteristics of the population-set. More precisely, the population-set's independent control actions statistically diversify themselves, thereby creating a smoother aggregate power consumption signal. The lack of coordination between TCL devices is observed to lead to significant load synchronization in aggregate power demand.

Load synchronization is again re-amplified after the demand response event. During this event, indoor air temperatures rise above their upper dead-band threshold, as shown in Figure 4.1a. After the reinstatement of power, all SSHPs begin cooling

their respective space within moments of one another. A significant spike in aggregate power consumption then occurs followed by repeated oscillations. The ringing observed in the hysteresis control architecture is precisely what our optimal control framework is trying to avoid.

The advanced hysteresis controller, first proposed by (Kuwada et al., 2020), eliminates much of this load synchronization through sparse network communication. Based on the aggregate power consumption subplot of Figure 4.2b, load synchronization is observed to decay quickly, with relatively minor sub-peaks observed in the FFT subplot of Figure 4.2c. Compared to its predecessor, the advanced hysteresis controller has a smoother aggregate power consumption signal.

As shown in Table 4.3, both of these low-level control frameworks have a similar efficiency, E_T , values. However, peak power consumption before and after the demand response vary drastically. Although not a faithful comparison, the PID control framework also experiences a large spike in peak power consumption after the demand response event. Of these three low-level controllers, load synchronization is observed to be a persistent issue.

In the United States the majority of air conditioning devices are SSHP and are controlled via the hysteresis controller of Equation (2.5). Naturally, the reader might ask “what might a optimal solution look like?” and “Can this optimal control framework mitigate load synchronization while maintaining indoor air temperatures?”. The simulation results of Section 4.2, show moderate improvements over their baseline study. When comparing the simulation results of the DD control framework to its low-level control counterparts, a smoother aggregate power consumption signal is observed in Figure 4.4b. Moreover, a smaller population size was simulated when observing the

optimal control frameworks. Unlike Figures 4.1b and 4.2b, Figure 4.4b shows smooth, albeit noisy, gradual changes in aggregate power consumption. Moreover the FFT subplot of Figure 4.4c shows a singular peak at $0Hz$ frequency with attenuated level of high frequency noise. The main difference between the advanced hysteresis and DD controller is the total time indoor air temperature spends above its upper dead-band threshold indicated by T_{dev} in Table 4.3. Comparatively, the DD controller mitigates all forms of load synchronization.

A significant aspect of this study was determining if sparse network communication performs similarly to an omniscient controller. Observing the aggregate power consumption subplots of Figures 4.4b and 4.6b, the DD has distinguishable higher frequency noise. However, as the number of simulated homes increases, this high-frequency noise is attenuated substantially. The DD shows smooth gradual changes whereas the CD framework tends to develop distinct aggregate power consumption bands. Another key difference between these two control frameworks is the inherent computational complexity. A drastic difference in simulation run-times are observed as the population size, N , increases.

Because Algorithms (5) and (6) compute control actions in a fundamentally different manner, comparing simulated run-times provides little meaning. The total number of simulated homes was chosen such that the CD, and to a lesser extent CC, control frameworks finished their simulation. These simulations represent a theoretical limit on what can reasonably be solved given standard desktop computer hardware. Both DD and DC are capable of computing any arbitrary population size because its computational complexity does not depend on the number of homes being simulated, N . Moreover, both decentralized control types are expected to benefit

from parallelization.

The DC and CC control frameworks have not only the highest efficiency, E_T , but also the smoothest aggregate power consumption subplots. In Table 4.3, DC and CC are the only frameworks whose P_f value is lower than its P_i value. By deduction, controlling the rate in which m changes, in turn, smooths aggregate power consumption, thereby benefiting the electric utility. Characteristically, both DC and CC delay rapid changes in aggregate demand, thereby granting an electric utility time to reorient themselves after a demand response event.

No noticeable difference is observed between the response of these two VSHP optimal control frameworks. Conclusively, this indicates that sparse network communication, in the sense of continuous control, performs on par with its omniscient counterpart.

Once Algorithm (1) begins updating thermal parameters, considerable deviations in each error signal is observed, as shown in Figures 4.5, 4.7, 4.9, and 4.11. We partially attribute this initial instability to the initialization process used to generate thermal parameters. During these early time steps, the matrix P_k of Algorithm (1) lacks sufficient knowledge to estimate new thermal parameter values. In a physical setting, the plant dynamics are considered to be unknown, or at the very least, capable of change. Usage of a true error signal is impractical because the actual thermal characteristics of a residential home are unknown. For this reason, we employ two update conditions to determine if newly estimated thermal parameter values result in the accurate portrayal of the plant dynamics.

The first update condition is mathematically expressed as $\kappa(P_k) \leq c_1$, where $\kappa(\cdot)$ is the condition number (see (Pyzara et al., 2011)), P_k is defined in Algorithm (1), and

c_1 is an empirically determined constant. An interpretation of the condition number is a measure of the sensitivity a solution given changes to the input. Exemplified by the initial instability of Figures 4.5, 4.7, 4.9, and 4.11, an ill-conditioned matrix P_k leads to poorly estimated thermal parameter values. Therefore, a conditional check is employed to determine how parameters are likely to change given varying input magnitudes. The second update condition, $\dot{\theta}_k \leq c_2$, checks if newly estimated parameters varies significantly between adjacent time-steps. This conditional check indicates periods of stability. In brief, the first conditional check determines the resistance to change, while the second condition looks at the volatility of newly estimated parameters.

Within the given simulated day, each optimal control framework shows a brief moment of instability followed by a slow convergence upon the simulated plant parameters. The initial instability is mainly attributed to the rank deficiency of the regressor and observer information. In the DC and CC control frameworks, a faithful convergence does not occur until after the demand response event. This lack of convergence is attributed to the phenomenon of persistence of excitation, see (Åström & Wittenmark, 2013). Once the demand response event forces its variation, then sufficient information is injected into the system, thereby allowing parameters converge. Unlike the two continuous control frameworks, both the DD and CD frameworks routinely introduce data rich state information via its discrete control. Such a system is observed to converge at a quicker pace and without the help of a demand response event. All four control frameworks are expected to converge given enough time and variability of control.

CHAPTER 5:

CONCLUSION

In this study, we build upon the well-established second-order equivalent thermal parameter model by introducing a recursive least squares algorithm used to predict the plants dynamical characteristics. Our main contribution in this study is the development of decentralized control frameworks which optimally schedule the control actions of a population of SSHPs and VSHPs alike. In addition, two centralized control frameworks are also explored. These centralized frameworks provide a convenient point of reference from which to compare the decentralized variants' simulation results. Based on the simulation results shown in Section 4, we show these optimal control frameworks outperforming their baseline study in terms of load synchronization and peak power reduction. The proposed optimal control frameworks eliminates undesirable load synchronization and peak aggregate power consumption through the proper scheduling of control actions while maintaining indoor air temperature. At the expense of moderately higher computational complexity, our MPC controller significantly improves peak power performance. Also, we observe that limited network communication provides similar performance attributes as an omniscient centralized control architecture. Despite the DD framework's discrete nature, a smooth aggregate power consumption is experienced.

An experimental verification of the results we provide in this thesis is left as

future work. Due to variations in each framework's optimization program, additional research is needed to test mixed population. Lastly, further research is needed to characterize the risk associated with cyber-security and network dropout.

REFERENCES

- Åström, K. J., & Wittenmark, B. (2013). *Adaptive control*. Courier Corporation.
- Barabási, A.-L., & Bonabeau, E. (2003). Scale-free networks. *Scientific american*, *288*(5), 60–69.
- Behboodi, S., Chassin, D. P., Djilali, N., & Crawford, C. (2018). Transactive control of fast-acting demand response based on thermostatic loads in real-time retail electricity markets. *Applied Energy*, *210*, 1310–1320.
- Callaway, D. S. (2009). Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, *50*(5), 1389–1400.
- Dunning, I., Huchette, J., & Lubin, M. (2017). Jump: A modeling language for mathematical optimization. *SIAM Review*, *59*(2), 295–320.
- Fuller, J. (2017). Residential module user’s guide.
 URL http://gridlab-d.shoutwiki.com/wiki/Residential_module_user's_guide
- Johnson, A. P. (2010). The history of the smart grid evolution at southern california edison. In *2010 Innovative Smart Grid Technologies (ISGT)*, (pp. 1–3). IEEE.

- Kazmi, H., Suykens, J., Balint, A., & Driesen, J. (2019). Multi-agent reinforcement learning for modeling and control of thermostatically controlled loads. *Applied energy*, 238, 1022–1035.
- Koch, S., Mathieu, J. L., & Callaway, D. S. (2011). Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. PSCC*, (pp. 1–7). Citeseer.
- Kuwada, J., Mehrpouyan, H., & Gardner, J. F. (2019). Design resilience of demand response systems utilizing locally communicating thermostatically controlled loads. In *ASME International Mechanical Engineering Congress and Exposition*, vol. 59438, (p. V006T06A045). American Society of Mechanical Engineers.
- Kuwada, J., Schwartz, R., & Gardner, J. F. (2020). Local communication in populations of thermostatically controlled loads. *ASME Journal of Engineering for Sustainable Buildings and Cities*, 1(3).
- Kwac, J., Flora, J., & Rajagopal, R. (2014). Household energy consumption segmentation using hourly data. *IEEE Transactions on Smart Grid*, 5(1), 420–430.
- Liu, M., & Shi, Y. (2015). Model predictive control of aggregated heterogeneous second-order thermostatically controlled loads for ancillary services. *IEEE transactions on power systems*, 31(3), 1963–1971.
- Malhame, R., & Chong, C.-Y. (1985). Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9), 854–860.

- Pyzara, A., Bylina, B., & Bylina, J. (2011). The influence of a matrix condition number on iterative methods' convergence. In *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, (pp. 459–464). IEEE.
- Ruelens, F., Claessens, B. J., Vandael, S., De Schutter, B., Babuška, R., & Belmans, R. (2016). Residential demand response of thermostatically controlled loads using batch reinforcement learning. *IEEE Transactions on Smart Grid*, 8(5), 2149–2159.
- Ruelens, F., Claessens, B. J., Vrancx, P., Spiessens, F., & Deconinck, G. (2019). Direct load control of thermostatically controlled loads based on sparse observations using deep reinforcement learning. *CSEE Journal of Power and Energy Systems*, 5(4), 423–432.
- Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., & Shelby, J. (2018). The national solar radiation data base (nsrdb). *Renewable and Sustainable Energy Reviews*, 89, 51–60.
- Tindemans, S. H., Trovato, V., & Strbac, G. (2015). Decentralized control of thermostatic loads for flexible demand response. *IEEE Transactions on Control Systems Technology*, 23(5), 1685–1700.
- Wang, J., Chen, X., Yu, K., Xie, J., Gan, L., & Xu, S. (2019). Optimal scheduling of air conditioning loads by aggregator under dynamic price. In *2019 IEEE 3rd Conference on Energy Internet and Energy System Integration (EI2)*, (pp. 1991–1995). IEEE.
- Zhang, W., Lian, J., Chang, C.-Y., & Kalsi, K. (2013). Aggregated modeling and

control of air conditioning loads for demand response. *IEEE transactions on power systems*, 28(4), 4655–4664.

Zhou, Y., Wang, C., Wu, J., Wang, J., Cheng, M., & Li, G. (2017). Optimal scheduling of aggregated thermostatically controlled loads with renewable generation in the intraday electricity market. *Applied energy*, 188, 456–465.

APPENDIX A:

DERIVATIONS

A.1 Second-order ETP model

Based on the thermal circuit shown in Figure 2.1, two resulting first order differential equations are formed, namely,

$$C_A \dot{T}_A = Q_A - \frac{1}{R_1} (T_A - T_o) - \frac{1}{R_2} (T_A - T_M), \quad (\text{A.1})$$

and,

$$C_M \dot{T}_M = Q_M + \frac{1}{R_2} (T_A - T_M). \quad (\text{A.2})$$

A second-order equivalent thermal parameter model is created from this thermal circuit by first rewriting Equation (A.1) in terms of T_M ,

$$T_M = R_2 C_A \dot{T}_A + \left(\frac{R_2}{R_1} + 1 \right) T_A - R_2 Q_A - \frac{R_2}{R_1} T_o. \quad (\text{A.3})$$

Next, compute the derivative of Equation (A.3) with respect to time,

$$\dot{T}_M = R_2 C_A \ddot{T}_A + \left(\frac{R_2}{R_1} + 1 \right) \dot{T}_A - R_2 \dot{Q}_A - \frac{R_2}{R_1} \dot{T}_o. \quad (\text{A.4})$$

Substitute Equations (A.3) and (A.4) into Equation (A.2) then simplify,

$$\begin{aligned} C_A C_M R_2 \ddot{T}_A + \left(C_M \left(\frac{R_2}{R_1} + 1 \right) + C_A \right) + \frac{1}{R_1} T_A = \\ Q_M + Q_A + C_M R_2 \dot{Q}_A + \frac{1}{R_1} T_o + C_M \frac{R_2}{R_1} \dot{T}_o \end{aligned} \quad (\text{A.5})$$

This concludes the derivation of the second-order ETP model.