

**Data-Driven Approaches to Demand Response:
Studies on Thermostatically Controlled Loads
and Electric Vehicles**

Submitted in partial fulfillment of the requirements for
the degree of
Doctor of Philosophy
in
Civil and Environmental Engineering

Emre Can Kara

B.S., Civil Engineering, Bogazici University, Turkey
M.S., Building Engineering, TU Delft, the Netherlands

Carnegie Mellon University
Pittsburgh, PA
December, 2014

©2014 Emre Can Kara. *Some rights reserved.* Except where indicated, this work is licensed under a Creative Commons Attribution 3.0 United States License. Please see <http://creativecommons.org/licenses/by/3.0/us/> for details.

The views and conclusions contained in this document are those of the author, and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government, or any other entity.

Keywords: demand response, state estimation, electric vehicles, thermostatically controlled loads.

Abstract

This dissertation presents a set of studies that use data collected on thermostatically controlled loads (TCLs) and electric vehicles (EVs) to *(i)* understand and improve their participation in demand response (DR) strategies, and *(ii)* quantify their capabilities as DR providers. First, we focus on TCLs and formulate a centralized control scenario in which a large, heterogeneous population of TCLs are controlled to provide ancillary services. We investigate the system-level benefits that such a control scenario can impart, and focus on evaluating modeling strategies that capture TCL population dynamics when disturbances to individual TCLs are considered. We then discuss the upstream communication requirements from the TCLs to the main controller and the characteristics of the underlying cyberinfrastructure, and argue that end users will prefer a strategy in which they do not have to share real-time information with the centralized controller. Using measurements obtained at a neighborhood-level load aggregation point, we develop state estimation techniques to replace the need for upstream information exchange between individual end use loads and the controller. Our results have important policy implications for appliance standards and the participation of loads in DR services. We show that the proposed aggregate TCL modeling strategy better captures the dynamics of an aggregate TCL population with no additional computational burden in comparison to state-of-the-art strategies. We also show that if individual thermal parameters of TCLs are available to the main controller, the controller can improve on its knowledge of the state of the TCL population us-

ing measurements obtained from a neighborhood-level load aggregation point.

Following our study of TCLs, we focus on centralized control of an aggregation of commercial EV charging stations. We investigate the benefits of managed EV charging to different stakeholders engaged in the reliable operation of the power grid. To do this, we leverage data collected from smart devices to overcome limiting assumptions commonly made in the literature regarding: *(i)* driving patterns, driver behavior and driver types; *(ii)* the scalability of a limited number of simulated vehicles to represent different load aggregation points in the power system with different customer characteristics; and *(iii)* the charging profile of EVs. As part of this study, we investigate the relationship between the EV infrastructure availability, EV load flexibility and benefits to stakeholders, a relationship which has implications for future improvements to DR programs. We show that managed EV charging can decrease the contribution of EV charging loads to the system peak load by approximately 40%, and reductions up to 24% in the monthly bills are possible for EV aggregations.

To my wife, Lindsay.

Acknowledgments

I am indebted to many people who gave me their support and guidance. This thesis would not have been possible without a very fortunate combination of people and geography. Foremost, I would like to thank my advisor, Dr. Mario Berges, who has always been a patient listener, an unlimited source of valuable ideas, a real friend, and an exceptional mentor throughout these four years. His inquisitive nature, hands-on approach, and ability to clearly see the bigger picture (while being bombarded by my questions) helped me hone my research abilities.

I am grateful to Dr. Gabriela Hug, who gladly agreed to collaborate with me in the last years of my Ph.D. She has been a true source of encouragement and inspiration when I needed them the most.

I would like to thank Dr. Bruce H. Krogh for his challenging questions and exciting discussions in my first two years at Carnegie Mellon. He has been a constant reminder of the value of perfection and the importance of loving what you do.

The other members of my thesis committee were also always there to provide guidance and help during my Ph.D. I am thankful to Dr. Zico Kolter and Dr. Scott Matthews for their time and insightful comments on my work.

I would also like to acknowledge Pacific Gas and Electric Company and ChargePoint LLP for providing the data used in parts of this dissertation, and Pennsylvania Infrastructure Technology Alliance, Hewlett Packard Labs, and the Office of Electricity Delivery and Energy Reliability of the U.S. Department of Energy for generously funding my work.

The former and current members of the InferLAB and members of the Advanced Infrastructure Systems group have been influential through fruitful discussions and contributed to my work significantly. Many thanks to In-Soo, Matineh, Leneve, Jingkun, Suman, and Adrian, among others.

I am grateful to all of my colleagues in the Grid Integration Group and the Demand Response Research Center in Lawrence Berkeley National Laboratory for hosting me for the final year of my research and making me feel a part of their group. In particular, I am grateful to Sila Kiliccote for her constant guidance and support, Jason MacDonald, Michaelangelo Tabone, Duncan Callaway, Salman Masyakeh and Douglas Black for valuable discussions and their contribution to my work.

Without the support of my friends, this journey would have been colorless. I am grateful to Batuhan and Ayse for their friendship and endless support, Onur and Gunay for being constant providers of laughter, Tugce and Loubna for being there whenever I needed a good listener, and Matineh for hours spent contemplating the true meaning of a Ph.D. I also would like to thank Semra, Ilke, Sinan and Cem for their friendship regardless of the time difference.

I am grateful to my wife Lindsay for giving me the motivation to carry on, and the Madisons for always being there whenever I needed the warmth of a home.

Lastly, I would like to thank my parents and my sister. Their love and support are invaluable to me. Thank you for enduring all of these years apart and supporting my decision to satisfy my curiosity.

Contents

1	Introduction	1
1.1	Goals and Scope	4
1.2	Contributions	6
1.3	Dissertation Overview	8
2	Quantifying Flexibility of Residential Thermostatically Controlled Loads for Demand Response	11
2.1	Dataset	13
2.2	Estimating Thermal Models	14
2.2.1	ARMA Model	16
2.2.2	Thermal Properties of Buildings	18
2.2.3	Validation	19
2.3	Demand Response Potential	24
2.3.1	Case Studies	26
2.4	Results	29
2.5	Chapter Conclusions	34
3	Modeling of Thermostatically Controlled Loads as Demand Response Agents	37

3.1	Overview of Existing Models	40
3.1.1	Modeling Individual TCLs	40
3.1.2	Modeling Aggregations of TCLs	41
3.2	Methodology	44
3.2.1	Dataset Description	45
3.2.2	Modeling Individual TCLs	45
3.2.3	Modeling Aggregations of TCLs	49
3.3	Results	55
3.3.1	Individual TCL Behavior	55
3.3.2	Plant-Model Mismatch	56
3.4	Chapter Conclusions	57

4 State Estimation of Thermostatically Controlled Loads for Demand

Response		61
4.1	Introduction	62
4.2	Related Work	64
4.3	Estimation of TCL States from Aggregate Power Measurements	66
4.3.1	Individual and Aggregate TCL Models	66
4.3.2	The Moving Horizon State Estimator	67
4.3.3	Kalman Filtering	68
4.4	Case Studies	70
4.4.1	Case Study I	71
4.4.2	Case Study II	72
4.5	Using Contextually Supervised Source Separation for State Estimation	75
4.6	Methodology	76
4.7	Dataset and Daily Load Simulation	79

4.8	Case Studies	80
4.8.1	Case Study I	83
4.8.2	Case Study II	90
4.9	Chapter Conclusions	93
5	Estimating the Benefits of Electric Vehicle Smart Charging at Non-Residential Locations	95
5.1	Dataset	99
5.1.1	Load Flexibility and Infrastructure Use	100
5.2	Smart Charging Strategy	105
5.3	Charging Infrastructure Owner’s Perspective	109
5.3.1	Problem Formulation	109
5.3.2	Case Study	112
5.3.3	Results	113
5.4	Distribution System Operator’s Perspective	119
5.4.1	Problem Formulation	120
5.4.2	Case Study	121
5.4.3	Results	121
5.5	Chapter Conclusions	123
6	Conclusions	125
6.1	Key findings	125
6.2	Discussion	129
6.3	Broader Impact	131
6.4	Future Research Topics	134

List of Figures

2.1	A daily snapshot of each data field used in this study for a sample household.	15
2.2	Average temperature setpoint for each house in the dataset, stratified by hour of day and by weekday-v-weekend. All setpoints are presented as deviations from the mean during HB3. Top panel shows average setpoints on weekdays; bottom panel shows weekends.	20
2.3	Distributions of estimated thermal properties of each house in the dataset.	21
2.4	Model prediction errors following setpoint changes of greater than $4^{\circ}F$. Load shed events (top) are positive set point changes, and load recovery events (bottom) are negative setpoint changes.	24
2.5	Event/control mechanism	25
2.6	The cumulative probability distribution of d_{shed} obtained by simulating 4405 HVAC instances, available in the dataset for CS-1, 2 and 3.	29

2.7	The behavior of the aggregate power consumption of the HVAC population for CS-2 and 3. The markers do not represent individual data points; they are placed to differentiate different DR curves from each other.	31
2.8	Energy savings w.r.t. no DR case for CS-2 and 3 between 9AM and 11PM.	32
2.9	Energy savings during event w.r.t. no DR case for CS-2 and 3.	33
2.10	Average power during event for CS-2 and 3.	34
3.1	State bin transition model adapted from [1] and [2]	43
3.2	Plant-model mismatch evaluation	54
3.3	Plant model mismatch using the proposed LTI strategy and the bLTI model adopted from [3] where f is the forcing parameter and each row corresponds to result of scenarios S1,S2 and S3	57
3.4	Training results for \mathbf{B}^{LTI} and \mathbf{B}^{bLTI} . The vertical color bar displays the mapping between the value of the each element in the matrices and the corresponding color used to represent it.	58
4.1	Estimation Model	70
4.2	Mean KL Divergence, $\hat{D}_{KL}(X_t \hat{X}_t)$, for Different Time Horizons, T	72
4.3	State Vector Elements, $x_{i,t}$ vs. Simulation Duration T_{tot} , for f : 12.5%, 50% and 100%	73
4.4	Mean KL Divergence, $\hat{D}_{KL}(X_t \hat{X}_t)$, for Each Forcing Value, f	75
4.5	CSSE Model	80
4.6	Comparison of CSSE performance for varying forcing parameters f (Scenario a)	88

4.7	Comparison of CSSE performance for varying μ parameters f (Scenario b)	89
4.8	Comparison of CSSE performance under the parameters given in case study II.	92
5.1	Centroids of zip code regions forming the VAPs	102
5.2	Number of sessions per month	103
5.3	Aggregate EV load shapes for January and December for all the VAPs	104
5.4	Average number of sessions per unique EVSE per day	104
5.5	The variation in load flexibility	105
5.6	Distributions of time-of-plug-in to EVSEs and time-of-plug-out from EVSEs.	105
5.7	Monthly bills calculated with E-19. The left bar for each month shows the current bill, and the right bar shows the optimized bill.	113
5.8	Decomposition of monthly bills to energy and demand charges. In each figure, the left bar shows the current charges, and the right bar shows the optimized charges for each month.	114
5.9	Charging-session-normalized Demand and Energy Charge Reductions	117
5.10	Demand Charge Reduction by Session Size	118
5.11	Distribution of percent peak shed for all the VAPs	122
5.12	Total energy moved outside of system peak period	123

List of Tables

2.1	Data field characteristics	14
2.2	Characteristics of different scenarios and case studies	28
3.1	Dataset used for HRUs [4]	45
3.2	Hyper-Parameters of T_i^{ON} and T_i^{OFF}	47
3.3	HRU parameters obtained from [5]	49
3.4	Individual simulation parameters for HRUs i	49
3.5	Different scenarios with varying $p_{d,i}$, h and f	55
3.6	Estimated parameters for HRU populations with a hundred thousand units: P1 and P2	55
4.1	Simulation Parameters for State Estimation	74
4.2	Function selection for case studies	81
4.3	Individual TCL parameters obtained from [5] for case study I	83
4.4	Input parameters for case studies I and II.	84
4.5	Individual TCL simulation parameters for case study II	90
5.1	VAPs used in this study	101
5.2	E-19 rate structure [6]	112
5.3	Average results based on summer and winter month rates in E-19	115

Chapter 1

Introduction

In a traditional power system architecture, generation is dispatched to match the load. Power system frequency regulation, voltage control and other key ancillary services are provided by controlling large-scale generators, transformers and other bulk-power devices [7]. This approach of matching generation to meet the load reflects an asymmetry in the roles played by supply and demand: small, individual loads see the power system as an *infinite bus*, whereas large-scale generators and devices see aggregate loads characterized by relatively slow-varying average values with additional light *noise* caused by thousands of independently varying small loads. In other words, traditional load models rely on the law of large numbers. Matching load to generation has been identified as one of the major challenges in the power grid of the future [8].

The task of maintaining power reliability and regulating frequency and voltage typically falls upon fossil fuel power plants, which remain idle for the majority of the time and then react during short periods of peak demand to provide power. For example, a study of New York and New England showed that for more than 88% of the time, the region was using less than 30% of its generation capacity [9].

In addition, as more consuming sectors grow dependent on digital devices, consumers are asking more from the power grid in terms of both capacity and reliability, and tolerance of power quality drops such as brownouts and blackouts has fallen [10]. The increasing demand for high-quality power is still met by inefficient infrastructure, and imbalances in demand and supply often yield blackouts and/or brownouts. Per year, U.S. customers are likely to experience 1.5 to 2 power interruptions and between 2 and 8 hours without power [9].

Furthermore, twenty-nine states in the U.S. have adopted renewable portfolio standards that require a certain percentage of their electricity generation to come from renewable generation sources [11, 12]. The intermittent, uncontrollable, and uncertain nature of renewable energy resources, which are increasingly being integrated into the power grid, requires additional reserves to keep the supply and demand balanced [10]. As an alternative to expanding the fossil-fuel based traditional resources, improving *demand response strategies* have been garnering interest in the research community, as studies have shown that they are capable of decreasing or eliminating the need for additional capacity [13, 11, 14, 15, 16, 17].

In a recent report, the U.S. Federal Regulatory Commission defined demand response (DR) as: "*Changes in electric usage by end use customers from their normal consumption patterns in response to changes in the price of electricity over time, or incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized*" [11]. Demand response programs have been implemented in the U.S. since the 1970s. Readers may be familiar with the use of utility-controlled thermostats that utilities offer to customers along with monetary incentives to reduce peak load on hot summer days. Although FERC's report does not take into account autonomous, decentralized control strate-

gies, it shows that demand response mechanisms could reduce the forecasted peak demand for 2017 by 14% [11]. However, they bring additional complexity to the operation of the power grid. These challenges stem from the ambiguity in individual end use loads' DR capabilities, the end user's willingness to participate in DR programs, and the type of DR services they are expected to provide.

Currently, most DR strategies are called on a day-ahead or hour-ahead basis, and often a human is involved in the control loop. Moreover, these strategies are based on an open-loop control approach, where a signal initiates a response and a building manager or energy management system takes preprogrammed control actions without providing any feedback to the entity that generated the initial signal [18]. Emerging DR strategies require moving from day-ahead participation to sub-hourly participation (i.e. ancillary services markets). This requires an understanding of: DR capabilities of the end use loads; end users' utility during DR; and which services are most suited to certain types of end use loads. Furthermore, it requires automated, fast response from end use loads and a close to real-time communication infrastructure [19].

Recent advances in technology have made it possible to embed computing devices in various entities connected to the power grid. Such devices open the door for observing and controlling large-scale phenomena with high temporal resolution, and make it possible to monitor and control different resources distributed across large areas and at varying time scales. For example, smart meters and sub-metering devices in buildings [20, 21] provide insight into consumption patterns; home automation networks make it possible to manage building energy consumption; and EV charging stations [22] provide information on the charging behavior of users, and allow for managed EV charging. We envision that data collected through these

devices can help us understand the end user's expectations from the power grid and develop strategies that support the reliable operation of the power grid without disturbance to end users.

The vision of this dissertation is *to develop demand response strategies that make use of the inherent flexibility of end use loads while ensuring that these strategies are non-disruptive to end users*. We envision leveraging data obtained from end use loads through sensors to propose improvements to the existing DR services and to quantify the benefits of DR services provided by these loads.

1.1 Goals and Scope

Investigating the possible benefits of all types of end use loads through all types of demand response programs would have limited the depth of this research. Therefore, in this dissertation, we specifically focus on direct load control approaches, where the response from individual loads is automated. To motivate the use of direct load control, we emphasize the need for an autonomous reaction in load consumption to a control signal as opposed to a behavioral change on the customer's end. Without automation, moving towards *fast* and *predictable* demand-side control while ensuring the reliability of the power grid is unrealistic and inefficient [11, 12, 23]. Responses to pricing signals are often unpredictable when the decision is left to the individual customer [24]. Programs that rely solely on customer response are more suitable for providing *slow* DR (i.e. day-ahead DR) services in which notifications regarding a DR event is sent to the DR providers a day before the event starts. By their nature, direct load control programs can provide faster and more predictable responses because individual loads are often automated and controlled directly based on grid instability or price information [24].

To keep the scope of this dissertation tractable and well defined, we will focus on two types of loads: thermostatically controlled loads (TCLs) and electric vehicle (EV) charging loads. TCLs constitute a significant portion (more than 50%) of total U.S. electricity consumption, and are widely available in most U.S. households [25]. It is possible to toggle them ON/OFF or control their power usage without influencing their end use function, and they can be modeled using heat transfer equations [13, 26]. All of these properties make TCLs favorable for a dynamic control mechanism capable of providing a faster response via DLC. EV loads are similarly favorable. First, EV charging loads are expected to increase significantly due to their role in achieving various states' greenhouse gas emission reduction targets [27, 28, 29]. Furthermore, the inherent flexibility in EV charging loads makes them favorable for providing DR services to the grid. Often, EVs stay plugged to the electric vehicle supply equipment for longer than what is required for a full state of charge. We believe that when this temporal flexibility is optimally managed for each EV by smart charging, DLC can provide system-level benefits such as peak-load reduction.

In this dissertation, our first goal is to investigate the system-level properties that can be influenced through the centralized control of a large, heterogeneous and smart collection of thermostatically controlled loads, where we use the term *smart* to denote their ability to react to measurements or signals. Recent work has shown the promise of this vision when incorporating large loads such as large HVAC systems. In contrast to these, we focus on aggregations of much smaller loads found in most residential buildings such as refrigerators, water/space heaters and small HVAC units. Specifically, we focus on understanding modeling strategies that capture the dynamics of TCL populations when disturbances to individual TCLs

are considered.

Second, we aim to understand the upstream communication requirements from the TCLs to the main controller and the characteristics of the underlying cyberinfrastructure. We believe that privacy plays an important role in end users' willingness to participate in a DR program; therefore, the upstream information transfer should be minimized while ensuring that the controller infers the necessary information.

Finally, we aim to quantify the benefits that centralized control of an aggregation of commercial EV charging stations can offer to different stakeholders engaged in the reliable operation of the power grid. Until now, most of the work concerning managed charging of EVs made limiting assumptions regarding: *(i)* driving patterns, driver behavior and driver types; *(ii)* the scalability of a limited number of simulated vehicles to represent different load aggregation points in the power system with different customer characteristics; and *(iii)* the charging profile of EVs. We believe data collected from smart devices can help to overcome these limitations when quantifying the benefits of centralized smart charging of EV aggregations.

1.2 Contributions

In this dissertation, we make the following contributions:

We test the following hypothesis: **a Markov Decision Processes (MDP) based model of a heterogeneous TCL population decreases the plant-model mismatch in comparison to a state-of-the-art linear-time-invariant model when disturbances to individual TCL operations are considered.**

Capturing the dynamics of thermostatically controlled load (TCL) aggregations, the end user’s interactions with individual TCLs and disturbances to their operation is key when evaluating their promises as demand response providers. Until now, researchers evaluating TCL aggregations used simulations of individual TCLs that made important assumptions, especially regarding the nature of the thermal characteristics assigned to the loads. We show that a data-driven modeling strategy to simulate individual TCLs—specifically, household refrigeration units (HRUs)—allows us to relax some of these assumptions and evaluate the validity of the approaches proposed to date. Building on the individual TCL model, an aggregate model is created using a bottom-up approach inspired by an MDP based strategy that results in a more flexible system. The performances of the proposed aggregate model and a state-of-the art LTI model are measured by quantifying the plant-model mismatch using root mean squared error.

We quantify the state estimation accuracy for tracking an aggregate TCL population from power measurements obtained at a neighborhood level using the proposed MDP based modeling strategy.

The quality and effectiveness of grid level services provided by the centralized control of TCLs depend highly on the communication requirements and the underlying cyberinfrastructure characteristics. Specifically, ensuring end user comfort while providing real-time demand response services depends on the availability of the information provided from the thermostatically controlled loads to the main controller regarding their operating statuses and internal temperatures. In an effort to replace the need for an upstream communication platform carrying information from appliances to the main controller in real-time, we study the accuracy with which inferences about these load-level properties (operating status and internal tempera-

ture) can be made from aggregate power measurements under different measurement assumptions. We use absolute and relative deviation metrics to quantify the performance of the proposed state estimation techniques.

We quantify the benefits of managed EV charging to distribution system operators and EV charging infrastructure owners by leveraging an extensive dataset of more than 500,000 commercial EV charging sessions.

The increase in the number of EVs in many states in the U.S. is expected to yield to operational challenges when these loads are considered inflexible. However, when considered flexible, they bring opportunities to support the reliable operation of the grid. Although many studies have attempted to quantify the benefits that managed charging of EVs can offer to distribution system operators and charging infrastructure owners, these studies make significant assumptions regarding trip and customer characteristics, number of vehicles and mobility patterns and the individual charging profile of the EVs. We use an extensive dataset collected from commercial EV charging stations to overcome these assumptions and to accurately quantify the benefits of managed EV charging to distribution system operators and the charging infrastructure owners while ensuring no hindrance to end users' comfort.

1.3 Dissertation Overview

In Chapter 2¹, we present a study that motivates the use of a unique and detailed dataset to estimate and validate models of demand response from TCLs—specifically, HVAC units—and quantifies the extent to which a population of TCLs

¹This chapter is based on joint work with Michaelangelo Tabone as published in [30].

can provide DR with preprogrammed setpoint changes. We use measured temperature setpoints, internal temperatures, compressor cycling ratio and metered energy data collected from over 4200 homes in Texas during the summer of 2012. Using autoregressive moving average (ARMA) models for individual households, we investigate the instantaneous power shed, the duration of the power shed, steady state energy savings and total energy savings. Specifically, we provide insight into the dependency of household DR availability to the temperature setpoint schedule, outdoor air temperature and time of day.

In Chapter 3², we shift our attention to the aggregate models of TCLs; leveraging physics based individual TCL models. The motivation behind this is threefold: *(i)* modeling individual TCLs becomes intractable as the number of TCLs increases, *(ii)* individual models of TCLs are more likely to invite privacy concerns from the end users, and *(iii)* the law of large numbers helps when modeling the aggregate behavior of TCLs. This chapter is dedicated to developing an individual TCL model, where thermal properties are obtained from distributions that capture the disturbances to the individual loads, and an MDP based model of a heterogenous TCL population. We then quantify the plant-model mismatch and evaluate the proposed strategy with a realistic simulator.

In Chapter 4³, building on the proposed aggregate model, we focus our attention on the communication requirements and the underlying cyberinfrastructure characteristics of a TCL population. Specifically, we aim to understand whether it is necessary for the TCLs to send their thermal properties to the main controller—a one-time information transfer—and whether the TCLs must also send data on their states to the main controller in real-time. To do this, we study the accuracy with

²This chapter is based on [31, 32, 33].

³This chapter is partially based on [2, 33].

which the state of the TCL population can be obtained from aggregate power measurements under different assumptions. We first introduce a moving horizon mean squared error state estimator with constraints as an alternative to a Kalman filter approach, which assumes a linear model without constraints. We develop this estimator assuming aggregate power measurements of individual TCLs are available. The results show that some improvement is possible for scenarios when loads are expected to be toggled frequently. Then, we focus on contextually supervised source separation techniques to elaborate further on the state estimation problem and propose a constrained optimization-based state estimation technique when aggregate power measurements are obtained at a distribution feeder serving a neighborhood of over 251 households.

In Chapter 5⁴, we use data collected from over 2000 commercial electric vehicle supply equipments (EVSEs) located throughout 190 zip code regions in Northern California spanning one year. To the best of our knowledge, this is the first study that uses such an extensive dataset on EV charging. First, we analyze over 580,000 charging sessions to investigate the trends in load flexibility and infrastructure use in the dataset. We shed light on the benefits of EV aggregations operated under a single owner, where a time-of-use pricing schema is used to study the benefits to the owner when EV aggregations shift load from high cost periods to lower cost periods. Then, following a similar aggregation strategy, we show that EV aggregations can significantly decrease their contribution to the system peak load.

Finally, in Chapter 6, we discuss the key findings of this dissertation. We discuss the results of each chapter and comment on the broader impacts of our findings. We conclude by describing future research directions.

⁴This chapter is based on [34].

Chapter 2

Quantifying Flexibility of Residential Thermostatically Controlled Loads for Demand Response

In recent years, thermostatically controlled loads (TCLs) such as heating, ventilation, and air conditioning (HVAC) systems, refrigerators, and water heaters have been garnering interest in the research community as DR resources. This is partly because of their wide availability and partly because of the thermal storage that they possess; this storage allows them to be turned on and off for an undetermined amount of time without affecting the performance of the thermostatic control.

Direct load control (DLC) of a population of residential TCLs has been studied by various researchers [31, 1, 35, 36, 37, 38, 39]. A linear time-invariant representation and a Markov decision process (MDP)-based representation of a heterogeneous

This chapter is based on joint work with Michaelangelo Tabone as published in [30]. This chapter is authored or co-authored by an employee, contractor or affiliate of the United States government. The United States Government retains a nonexclusive, royalty-free right to publish or reproduce this chapter, or to allow others to do so, for Government purposes only.

TCL population using state-bin transition models is given in [5] and in [31], respectively. A priority stack-based algorithm to provide ancillary services is presented by [38]. The authors of [40] use household refrigeration units that are modified with additional thermal storage to provide peak shaving.

In addition, the quantification of TCLs' capability to provide DR has also been investigated by various researchers through bottom-up simulation methods, where simulation model parameters are based on educated guesses [38, 41]. Studies investigating the impacts of DR on the power grid level include, but are not limited to, [42, 26, 43, 44].

However, in existing studies, quantification of the resources are based on bottom-up simulations using plausible but not rigorously identified parameters. To the best of our knowledge, [45] is the only work that uses real electricity consumption data to estimate models that can then be used to compute DR potential. However, the dataset used in [45] consists of whole-building smart meter data, so a great deal of the modeling effort (and uncertainty) is due to the challenge of disaggregating HVAC demand from whole-building demand.

This study's central contributions are (i) a method to use direct HVAC measurements to estimate DR potential, and (ii) initial estimates of DR potential from a large aggregation of residential loads using this method. We use a unique dataset in which HVAC consumption, temperature setpoint and measured indoor temperature were recorded for over 4200 households. This allows us to sidestep the issue of disaggregation encountered in [45]. Using autoregressive moving average (ARMA) models trained based on historical data from individual households, we investigate the impacts of varying exogenous parameters on flexibility metrics such as instantaneous power shed, the load shed duration, steady-state energy savings and total

energy savings. We investigate the dependency of household DR availability to the temperature setpoint schedule, outdoor air temperature and time of day.

We find that the constant exogenous input assumption made commonly in the literature (e.g. [31, 1]) yields to a significant over-estimation of the number of loads that can provide a requested shed duration under certain cases. For the cases with a larger setpoint adjustment, the steady-state energy savings are higher for both weekday and weekend profiles.

The chapter is organized as follows: Section 2.1 describes the dataset used in this study. Section 2.2 introduces the ARMA model built on [46] and evaluates the performance of this model for a single household. Section 2.3 uses the model proposed in Section 2.2 to estimate the DR flexibility, and discusses the sensitivity of these estimations to exogenous inputs. Section 2.4 presents the results of the DR flexibility estimation, and finally, conclusions are drawn and future work is suggested in Section 2.5.

2.1 Dataset

The dataset used in this study is provided by Landis&Gyr, and it is collected from 4297 households located in Texas. The data include temperature setpoint, indoor temperature, compressor cycling ratio and energy consumption measurements sampled every 5 minutes during the summer of 2012. The compressor cycling ratio is defined as the ratio of the time that the HVAC compressor is ON within a sampling period to the sampling period. In addition, hourly weather data obtained from weather stations in the Texas area closest to the households are used to capture the external temperature. The energy consumption measurements are converted into average power measurements assuming constant power use within each sampling

Table 2.1: Data field characteristics

Data Field	Resolution	Sampling Rate
Temperature setpoint, $T_{set}(t)$	1 [$^{\circ}F$]	5 mins
Internal Temperature, $T_{int}(t)$	1 [$^{\circ}F$]	
Duty Ratio, $d(t)$	0.001	
Average Power, P_{avg}	0.01 [W]	
External Temperature, $T_{ext}(t)$	0.1 [$^{\circ}C$]	1 hour
Solar Insolation, $\phi_{sol}(t)$	0.01 W/m^2	

interval. A solar earth geometry model was used to find the solar insolation on a horizontal surface at each home’s approximate location at each time of day, assuming no cloud cover. Specific data fields used in this study along with their sampling rate and resolution information are given in Table 2.1. Data was only collected from homes with communicating thermostats, and thus are not a representative sample of all homes in the service territory. However, if a demand response program is to only control homes which have communicating thermostats (a very likely case), the sample does represent homes that are currently capable of participating.

A daily snapshot of each data field in Table 2.1 for a sample household is given in Figure 2.1.

2.2 Estimating Thermal Models

We fit thermal models to each building monitored in our dataset with the objective of simulating transient dynamics of internal temperature and power consumption during an DR event. Two important parameters of our model were apparent from the data without detailed statistical models: average temperature setpoint of the thermostat, and instantaneous HVAC power consumption. We used linear regression to fit an ARMA model to account for appliance efficiencies, thermal masses, thermal resistances, interior heat gains and solar gains.

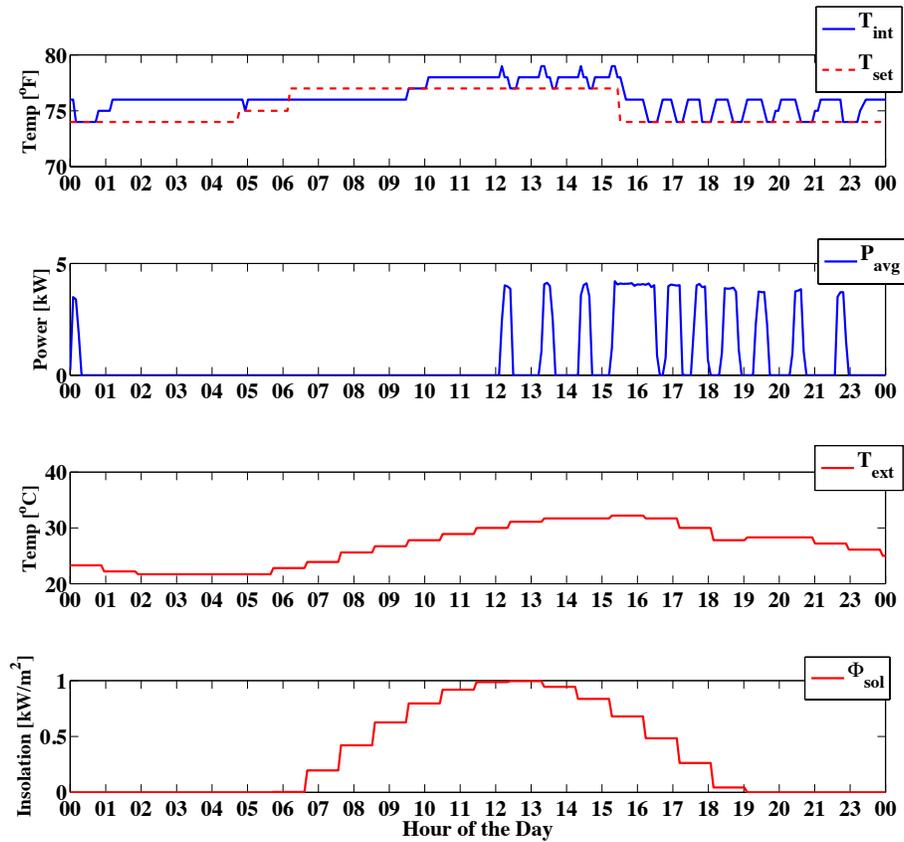


Figure 2.1: A daily snapshot of each data field used in this study for a sample household.

ARMA models allow predictions of internal temperature to be dependent on both coincident and prior (a.k.a. “lagged”) values of physical inputs such as cooling energy, outdoor temperature, and solar isolation. In the remainder of this section, we fully describe the ARMA model and detail how we constrained the parameters such that physical principles hold at steady state.

2.2.1 ARMA Model

The following equation shows an ARMA model for the thermal dynamics of a building, reproduced from [46]:

$$\begin{aligned} \sum_{i=0}^{N_{int}} a_{int}(i)T_{int}(t-i) &= a_{occ} + \sum_{i=0}^{N_{out}} a_{ext}(i)T_{ext}(t-i) + \\ &\sum_{i=0}^{N_{sol}} a_{sol}(i)\phi_{sol}(t-i) + \sum_{i=0}^{N_{aux}} a_{aux}(i)\phi_{aux}(t-i) + \epsilon(t) \end{aligned} \quad (2.1)$$

In (2.1), $T_{int}(t)$ is the room temperature of the house at time t , $T_{ext}(t)$ is the outdoor ambient temperature, $\phi_{aux}(t)$ is the average auxiliary power over the interval (from an air conditioner or a heater), and $\phi_{sol}(t)$ is solar insolation. The summation terms include time-lagged readings of each variable; N_x is the number of lagged readings included for the variable defined by the subscript x . Model coefficients are denoted as $a_x(j)$ where the subscript, x , indicates the corresponding variable and j indicates the timing of the lag (in number of readings prior to present). a_{occ} is an intercept term which is analogous to a constant internal heat gain resulting from occupants and devices. We applied two constraints to the ARMA model. The first constraint simply scales all coefficients such that $a_{int}(0) = 1$. The second constraint ensures that the steady state thermal properties are consistent and was originally presented in [46].

$$\bar{T}_{int}\bar{a}_{int} - \bar{T}_{ext}\bar{a}_{ext} = \bar{\phi}_{aux}\bar{a}_{aux} + \bar{\phi}_{sol}\bar{a}_{sol} + a_{occ} \quad (2.2)$$

Equation (2.2) shows the rearranged ARMA model from (2.1) where all variables are in steady-state; we define steady-state variables as $\bar{X} = X(t) \forall t$, and steady state coefficients as $\bar{a}_x = \sum_{i=0}^{N_x} a_x(i)$. Equation (2.2) is arranged such that all values

on the right-hand side represent heat flows into or out of the building, and values on the left represent indoor/outdoor temperatures. At steady state, heat transferred across the constant positive indoor/outdoor temperature differential ΔT should be equal and opposite to that transferred across its negative, $-\Delta T$, implying that the sum of the coefficient on internal temperatures are equal to the sum of coefficients on outdoor temperatures. We rewrite (2.2) as

$$L(\bar{T}_{int} - \bar{T}_{ext}) = \bar{\phi}_{aux}\bar{a}_{aux} + \bar{\phi}_{sol}\bar{a}_{sol} + a_{occ} \quad (2.3)$$

where L is defined as

$$L = \sum_{i=0}^{N_{int}} a_{int}(i) = \sum_{i=0}^{N_{out}} a_{ext}(i) \quad (2.4)$$

Rearranging this equality yields to the second constraint:

$$a_{int}(N_{int}) = \sum_{i=0}^{N_{out}} a_{ext}(i) - \sum_{i=0}^{N_{int}-1} a_{int}(i) \quad (2.5)$$

Substituting both constraints into (2.1) and solving for the latest internal temperature results in

$$\begin{aligned} T_{int}(t) - T_{int}(t - N_{int}) &= \sum_{j=1}^{N_{int}-1} a_{in,j}(T_{int}(t - j) - T_{int}(t - N_{int})) \\ &+ \sum_{j=0}^{N_{ext}} a_{out,j}(T_{ext}(t - j) - T_{int}(t - N_{int})) \\ &+ \sum_{j=0}^{N_{sol}} a_{sol,j}\phi_{sol}(t - j) \\ &+ \sum_{j=0}^{N_{aux}} a_{aux,j}\phi_{aux}(t - j) \\ &+ a_{occ} + \epsilon_t \end{aligned} \quad (2.6)$$

We fit the parameters shown in (2.6) for each household in the dataset using ordinary least square regression. We did not measure auxiliary heating/cooling energy directly; as a proxy, we substituted the average power consumption of the air conditioner (AC) over the interval. This substitution implicitly assumes that the coefficient of performance (COP) of the AC is constant during the study period, thus the power consumption of the AC is directly proportional to the cooling energy provided by the AC. In actuality, the COP decreases as the difference between indoor and outdoor temperature increases. We explore the effects of this assumption in the conclusions.

For the purposes of this chapter, the model order was chosen by visual inspection on a few houses: $N_{int} = 20$, $N_{ext} = 0$, $N_{aux} = 2$, and $N_{sol} = 0$. A more appropriate model selection process will be developed for future work. Twenty lagged internal temperatures were more than sufficient for most homes, as indicated by coefficients of long lags estimates to not be significantly different from zero. However, including superfluous lags did not affect the performance of the model and gave a conservative estimate of the model's complexity. Coefficients for lagged outdoor temperatures were difficult to identify and were left out of the model. Outdoor temperature was measured at roughly hourly intervals and then linearly interpolated; thus, there was not much variation at short lags. Including a few lagged values of auxiliary power was necessary, as there was often a noticeable delay between energy consumed by the HVAC and any response in room air temperature.

2.2.2 Thermal Properties of Buildings

Figures 2.2 and 2.3 show box plots of estimated properties of each building in the dataset. In these plots, boxes represent the interquartile range (25th to 75th percentile), and whiskers represent the 2.5th and 97.5th percentiles.

The **setpoints** of thermostats in our dataset varied frequently and regularly, presumably in an attempt to conserve energy by mitigating heat losses. Figure 2.2 presents distributions of the average setpoint of each thermostat at each hour ending compared to the average during the hour beginning at 3AM (HB 3); results are also stratified by weekday versus weekend. As shown, during the week, most households raised their setpoints by up to $4^\circ F$ during the afternoon as compared to the night. Fewer than 5% of households lowered their setpoints by more than $5^\circ F$, or raised it by more than $10^\circ F$. During the weekend, households were less likely to raise their setpoints during the afternoon.

Panel A of Figure 2.3 shows the distribution of estimated rated power use for all HVACs in the dataset, which is defined as the average power used by HVACs during intervals when the compressor cycling ratio is 1. Most HVACs consume between 2 and 4 kW of electricity when running at full load, though estimates vary widely, from 0.5 to 6 kW.

2.2.3 Validation

Panel B of Figure 2.3 shows the distribution of **steady-state heat transfer coefficients**, h_{ss} , estimated for each home in the dataset. h_{ss} represents the amount of additional power required by the air conditioner to increase the steady state indoor/outdoor temperature differential by $1^\circ F$. Equation (2.7) shows our formula for calculating h_{ss} , which is derived by rearranging (2.3):

$$h_{ss} = \frac{L}{\bar{a}_{aux}} = \frac{\sum_i a_{int}(i)}{\sum_i a_{aux}(i)} \quad (2.7)$$

Our definition of h_{ss} deviates from convention because we reference it to power consumed by the HVAC rather than cooling energy produced by the HVAC (which

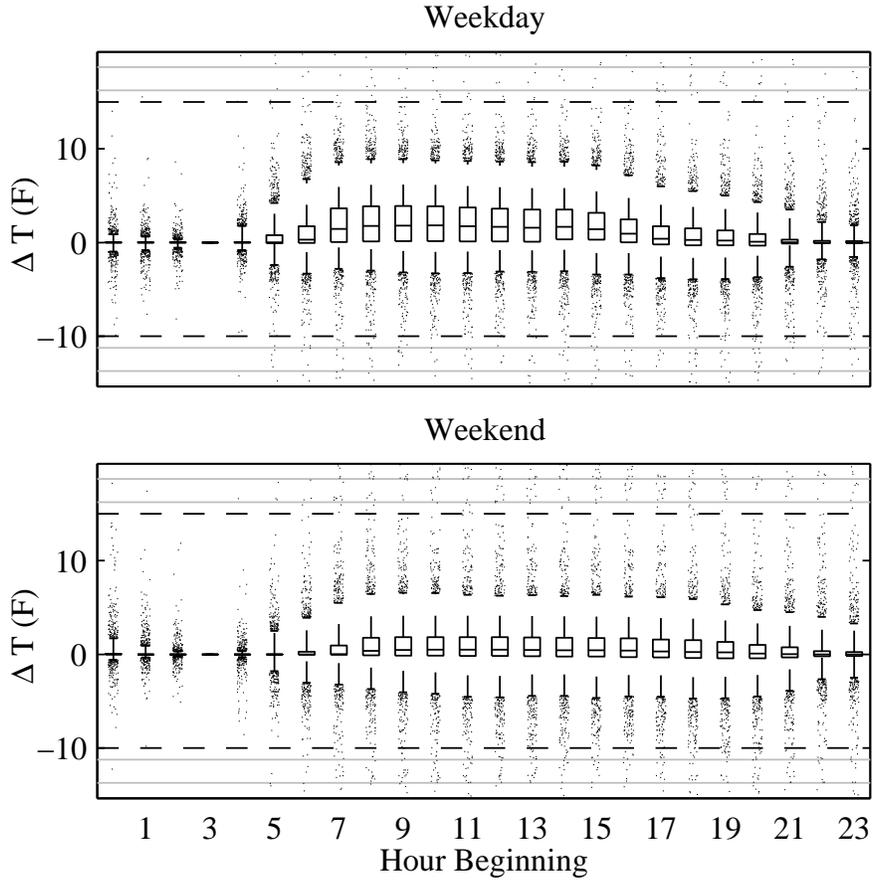


Figure 2.2: Average temperature setpoint for each house in the dataset, stratified by hour of day and by weekday-v-weekend. All setpoints are presented as deviations from the mean during HB3. Top panel shows average setpoints on weekdays; bottom panel shows weekends.

we do not measure). Thus, our coefficient accounts for (1) the thermal resistance of the building shell material, (2) the building shell area, and (3) the COP of the air conditioner. Most of these values are around $100 \text{ W}/^\circ\text{F}$; however, they vary from less than 0 (obviously erroneous estimates) to $800 \text{ W}/^\circ\text{F}$.

Panel C of Figure 2.3 shows the distribution of **effective occupant heat gain**, Q_{occ} , defined in (2.8) as the steady state, average HVAC power required to exactly

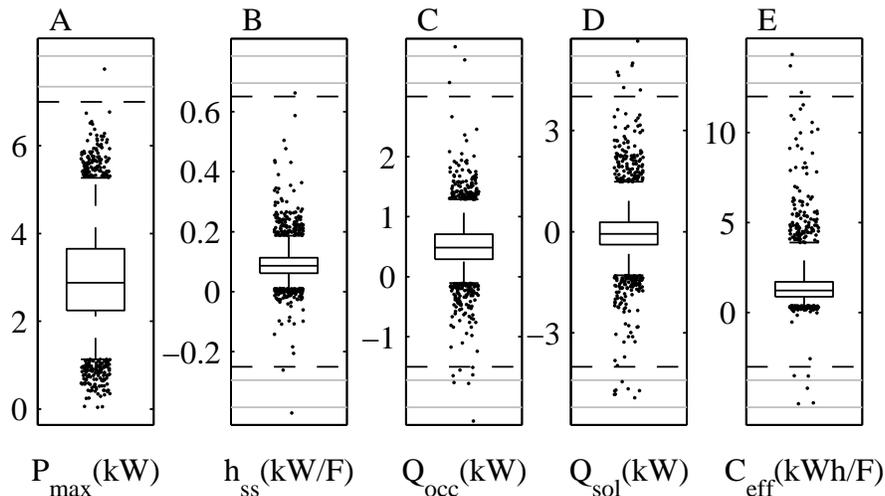


Figure 2.3: Distributions of estimated thermal properties of each house in the dataset.

offset the occupant heat gain:

$$Q_{occ} = \frac{a_{occ}}{\bar{a}_{aux}} = \frac{a_{occ}}{\sum_i a_{aux}(i)} \quad (2.8)$$

Dividing Q_{occ} by h_{ss} gives the indoor/outdoor temperature differential sustained by only the occupant heat gain and no HVAC use.

Panel D of Figure 2.3 shows the distribution of **maximum solar heat gain**, Q_{sol} , which is defined in (2.9) as the steady state average HVAC power required to offset solar heating gains at $1000W/m^2$:

$$Q_{sol} = \frac{1000 \sum_i a_{sol}(i)}{\sum_i a_{aux}(i)} \quad (2.9)$$

As shown, distribution of these gains is centered at 0, meaning that the model parameters indicate that 50% of buildings derive *cooling energy* from solar insolation (or more specifically, from our solar-earth geometry model-derived estimates of clear-sky insolation). This result clearly indicates that these model parameters are not

capturing the desired effect. This is likely a result of (1) shortcomings of using solar-earth geometry model data rather than solar observations (which would include cloud cover), (2) colinearity between the solar diurnal cycle and occupancy and (3) colinearity between the solar diurnal cycle and outside air temperature. We discuss these effects further in the conclusions.

Panel E of Figure 2.3 shows the distribution of **effective thermal capacity**, C_{eff} , which is defined in [46] as the amount of energy released by the building when reducing from a sustained internal temperature above steady state to the steady state. The value of C_{eff} is independent of all external variables as long as they are constant over time. We estimated this value by simulating the energy released by the building when returning from a steady temperature $1^\circ F$ above steady state. The capacity of buildings in the dataset is typically between 1 and 2 $kWh/^\circ F$. However, some buildings exhibit very large heat capacities.

We validate the model by predicting the response of internal temperature to large changes in setpoints that exist in the dataset. This validation metric is chosen to mimic the transient dynamics of an actual DR event which is often initiated via changing the setpoints of thermostats located in the households. Internal temperature was predicted based on knowledge of the internal temperature prior to the change (but not following), and the HVAC power consumption, outdoor temperature, and solar insolation throughout the time period.

In order to be considered for validation, setpoint changes must be $4^\circ F$ or larger and the setpoint must be constant for the two hours before and four hours after the change. The motivation behind this is to capture the dynamics of a population that is close to steady-state, while ensuring that there are enough validation cases

mimicking DR events that last for up to 4 hours after the setpoint change. We refer to a setpoint increase as a “load shed” and a setpoint decrease as a “load recovery,” analogous to the beginning and end of an DR event. There were 54,069 load sheds and 34,534 load recoveries with complete data in the dataset, and all were used for validation.

Figure 2.4 shows box plots of the model prediction errors at 1, 2, 3, 4, 5, and 6 hours following a setpoint change. The longer the time duration since the set point change, the greater the magnitude of the prediction errors. This result is expected because the setpoint change also marks the latest observations of actual internal temperature used to inform the prediction; predictions made six hours ahead of the last piece of information are more uncertain than those made one hour ahead.

Median prediction errors suggest that our model may be overestimating the effect of each building’s thermal capacity. For load shedding events, the model tends to under-predict internal temperature directly following the change, signifying that the building is increasing in temperature more quickly than the model is predicting, and thus has less thermal capacity than predicted. However, for load recovery events, the model appears to be unbiased for the entire duration following the setpoint change.

The over-prediction of thermal capacity could be a result of an under-prediction of the effects of air conditioning and thermal gains on the internal temperature. In the conclusions, we suggest that this could be an effect of the low resolution of the internal temperature observations ($1^{\circ}F$), as the rounding errors will be correlated with the effects of the cycling air conditioner.

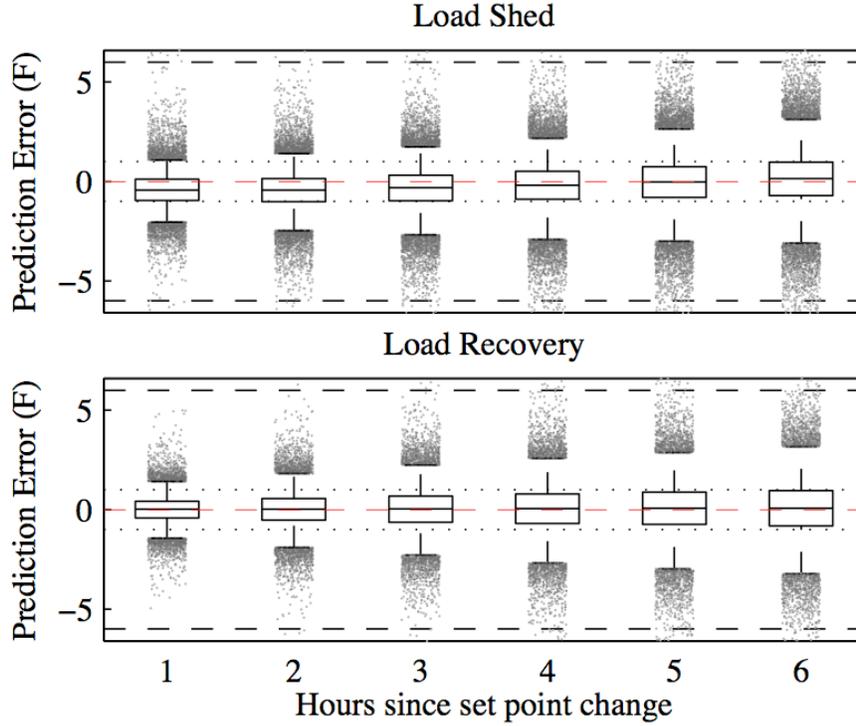


Figure 2.4: Model prediction errors following setpoint changes of greater than $4^{\circ}F$. Load shed events (top) are positive set point changes, and load recovery events (bottom) are negative setpoint changes.

2.3 Demand Response Potential

We use the ARMA model introduced in (2.6) to investigate the DR potential of an HVAC population under varying exogenous conditions. We simulate DR events by introducing a setpoint change in the thermostats, and we quantify the duration of the load shed $d_{i,t}^{shed}$ given a setpoint change ΔT_t^{set} under different conditions. Furthermore, we investigate the behavior of the aggregate power consumption of the HVAC population and estimate potential energy savings due to shifting the HVAC loads to much cooler periods of the day.

Formally, we define an DR event as follows: The internal setpoint T_{set} of the building is increased by a predetermined amount ΔT_t at time t . The duration of the

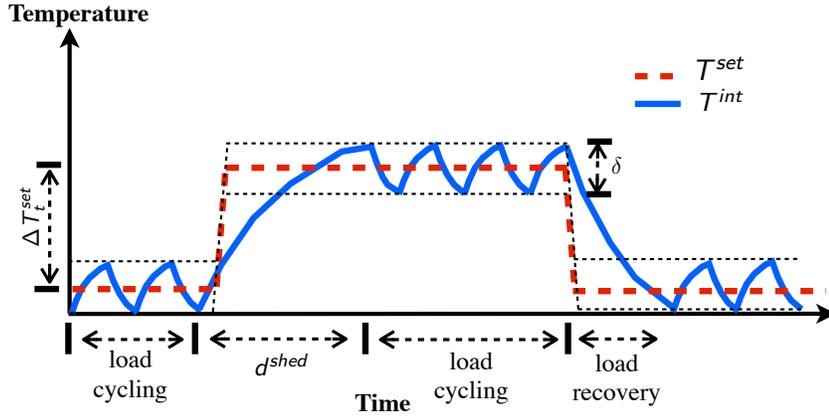


Figure 2.5: Event/control mechanism

load shed $d_{i,t}^{shed}$ for HVAC i at time t is defined as the time that it takes for $T_{i,t}^{int}$ to reach $T_{set} + \Delta T_t + \delta/2$, where δ is the thermostatic dead-band width. We simulate the behavior of each load as a 2-state load, whose thermal dynamics are governed by (2.6) and a thermostat. Specifically, we assume that the load turns OFF when the internal temperature hits the upper thermostatic bound (i.e. $T_{int} = T_{set} + \delta/2$) and vice versa. Figure 2.5 represents a typical load shed DR event considered in this study and depicts the parameters defined above.

As expected, time varying setpoints affect a building’s response to an DR event. If the temperature setpoint is maintaining a high differential between indoor and outdoor temperature, an DR event is expected to create a large instantaneous magnitude of power reduction, with a short duration (due to the high heat transfer over the differential). Conversely, if the temperature differential is small, the DR program should expect more modest instantaneous power reductions for a longer duration.

To capture this variation, we created various case studies using different setpoint profiles and exogenous input characteristics. We first investigate a simplistic scenario, where we assume all of the HVACs have the same temperature setpoint, the

external air temperature is constant, and the insolation is zero. We then select a day to use as a representative day for the exogenous inputs such as: external temperature and insolation. We use average hourly setpoint schedules for HVAC loads obtained from the weekday and weekend measurements to obtain distributions of the $d_{i,t}^{shed}$ under varying conditions. The next section introduces the case studies in detail and discusses the assumptions made.

2.3.1 Case Studies

For all the case studies in this chapter, we simulate a population of HVACs using the ARMA model given in (2.6). The simulation period and the thermostatic dead-band width is pre-determined and for all case studies t_{start} is 9AM, t_{end} is 11PM and δ is $1^\circ F$. At the beginning of each simulation, we assume the starting temperature for each HVAC, $T_{i,t_{start}}^{int}$, was randomly located within the thermostatic dead-band width. The HVAC statuses are also determined by observing the distribution of HVAC statuses at the starting hour from measured data based on the setpoint profile. During the simulation period, we assume that a single DR event happens at t_{event} , with a predetermined setpoint change $\Delta T_{t_{event}}^{set}$. Since the Electric Reliability Council of Texas' (ERCOT) DR programs have varying dispatch durations ranging from 1 to 4 hours, we simulate the behavior of the aggregate load population with varying DR event durations D_{DR} . Specifically, after a setpoint change of $\Delta T_{t_{event}}^{set}$ due to a DR event at time t_{event} , the setpoint of each HVAC is overwritten by $T_{i,t}^{set} + \Delta T_{t_{event}}^{set}$ for the duration of the DR event, D_{DR} . When the event is over, all HVACs go back to their scheduled setpoint profile.

Three sets of different case studies (CS-1, CS-2 and CS-3) are created to capture different exogenous input conditions. For each of these sets, we define 4 individual

scenarios (e.g. A-D for CS-1) with varying D_{DR} , t_{event} and $\Delta T_{t_{event}}^{set}$ parameters. As seen in Table 2.2, the first set of case studies CS-1 focuses on estimating the DR potential of an HVAC aggregation under constant exogenous inputs. For the second set, we create an average weekday setpoint profile for all the HVACs, and assume that each HVAC is scheduled based on that. Similarly for the third set, we create an average weekend setpoint profile to differentiate the varying setpoint profiles due to change in occupancy patterns between weekdays and weekends.

Table 2.2: Characteristics of different scenarios and case studies

Case Study	External Temp.	Setpoint Profile	Setpoint Change $\Delta T_{t_{event}}^{set}$	Scenarios	Event Start Hour, t_{event} , [hour of the day]	DR event Duration, D_{DR} , [hours]
CS-1	82°F, Constant	76°F, Constant	2°F	A	12	[1,2,3,4] hours
			4°F	B	15	
				C	12	
				D	15	
CS-2	Measured, Thu, 21/06/2012	Avg. Weekday Profile	2°F	A	12	[1,2,3,4] hours
			4°F	B	15	
				C	12	
				D	15	
CS-3	Measured, Sun, 17/06/2012	Avg. Weekend Profile	2°F	A	12	[1,2,3,4] hours
			4°F	B	15	
				C	12	
				D	15	

2.4 Results

Figure 2.6 shows the empirical cumulative distribution functions of $d_{i,t}^{shed}$ for the constant exogenous input case (CS-1), weekdays (CS-2) and weekends (CS-3). In

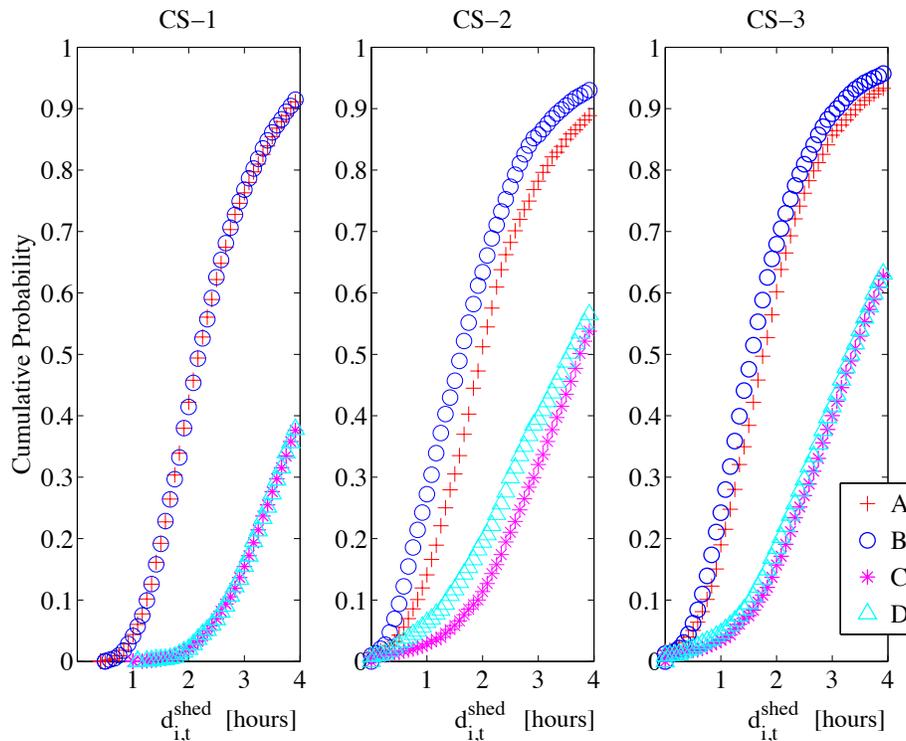


Figure 2.6: The cumulative probability distribution of d_{shed} obtained by simulating 4405 HVAC instances, available in the dataset for CS-1, 2 and 3.

all case studies, for any shed duration value d such that $d \leq 4$, it is possible to see that scenarios C and D have a lower $Pr(d_{i,t}^{shed} \leq d)$ value. This is expected because the setpoint change value for both C and D is $4^\circ F$, hence given a shed duration value, there are less loads available provide the requested shed at least for d long. For CS-1, since the exogenous inputs are assumed to be constant during the day, scenarios A, B, C and D yield to almost identical results. When the results for CS-1 is compared to CS-2 and 3, it is possible to see that constant exogenous input

assumption yields to a significant over-estimation of the number of loads that can provide a requested shed duration d for scenarios C and D. For scenarios A and B no significant difference is observed, which can be explained by the presence of less variation in the exogenous inputs during shorter shed periods that can be achieved by a lower setpoint change. For CS-2 and 3, it seems that scenario C has the lowest $Pr(d_{i,t}^{shed} \leq d)$ value among all 4 scenarios. This suggests that an DR event at 12PM has the highest chance of exceeding a shed duration of 4 hours for both weekdays and weekends.

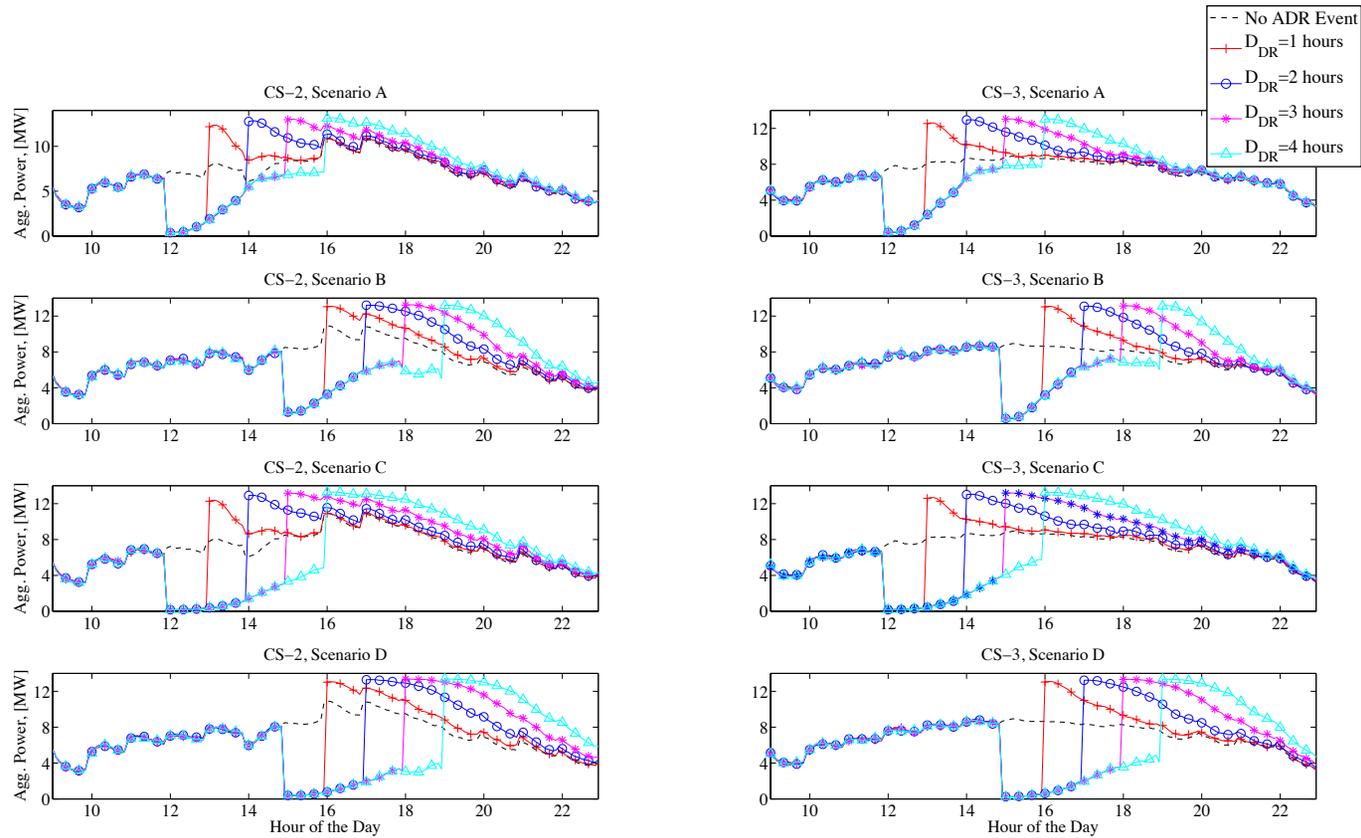


Figure 2.7: The behavior of the aggregate power consumption of the HVAC population for CS-2 and 3. The markers do not represent individual data points; they are placed to differentiate different DR curves from each other.

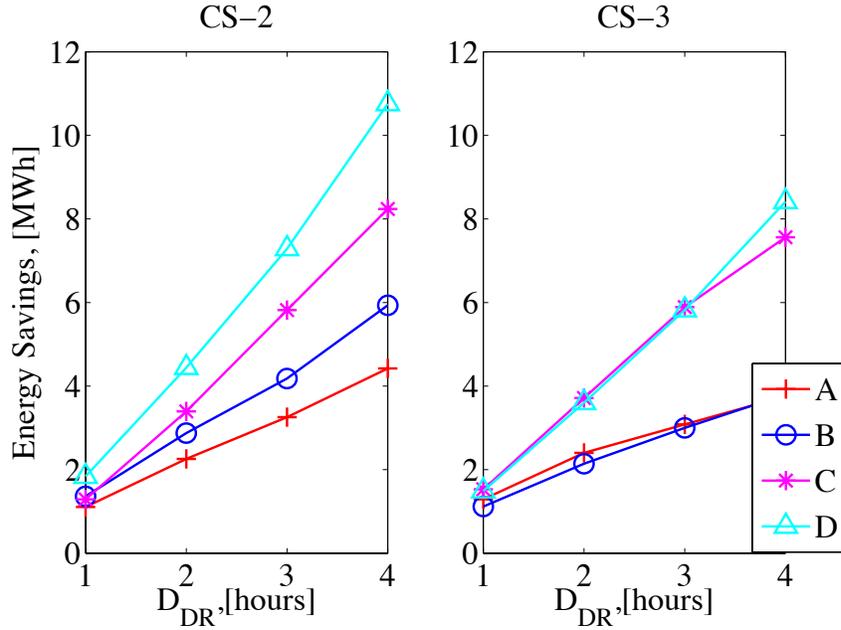


Figure 2.8: Energy savings w.r.t. no DR case for CS-2 and 3 between 9AM and 11PM.

Figure 2.7 shows the aggregate power time series obtained for each scenario under CS-2 and CS-3 against a no DR event case with identical input conditions. We choose to depict CS-2 and 3, since they represent a more realistic scenario. For all the scenarios with an DR event starting at 12PM, it is possible to see that the recovery peak value is increasing with the increasing D_{DR} . Since the maximum value of D_{DR} we investigated in this study is 4 hours, the benefits due to decreasing exterior temperature (if any) seems to be diminished. In contrast, for the scenarios with DR events scheduled at 3PM, a decrease in the rate of increase of the aggregate power consumption is often visible after 6PM, which helps to decrease the expected increase in the power consumption during recovery.

In Figures 2.8, 2.9, and 2.10, we depict the aggregate energy savings during the entire simulation period, energy savings during the event and the average power consumption during the DR event. We observe up to 11MWh energy savings in

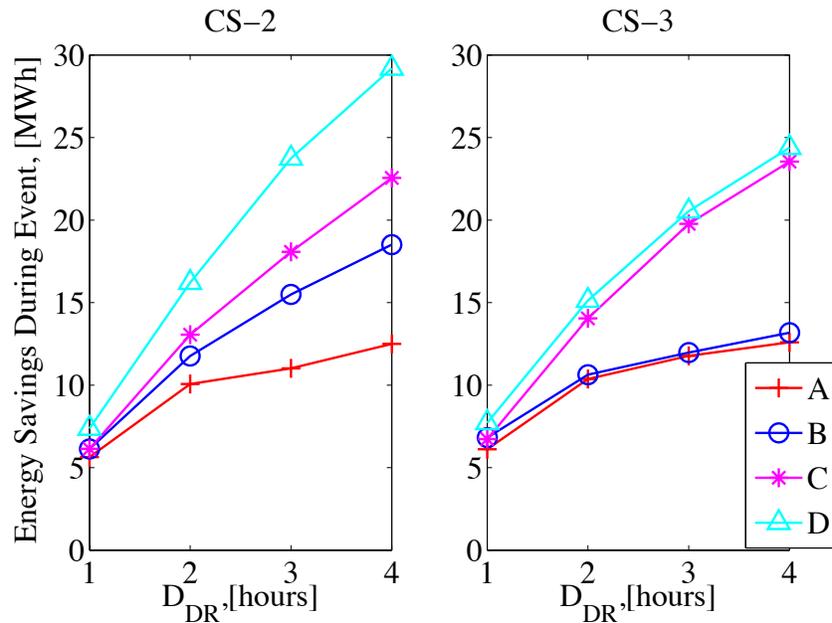


Figure 2.9: Energy savings during event w.r.t. no DR case for CS-2 and 3.

comparison to a no DR case between 9AM and 11PM, and up to 29MWh energy savings during the event period. It is possible to see that scenarios D and C can offer more overall energy savings than scenarios A and B for weekdays and weekends. It is also possible to see that a later DR event (3PM in this case) can provide additional benefits during weekdays, which can be explained by the expected increase in the occupancy in households after earlier events resulting in a lower recovery setpoint. However, on the weekends, there is no significant difference between earlier and later events in terms of overall energy savings. In addition, the rate of increase in the energy savings with increasing D_{DR} shed duration during the DR event diminishes with increasing D_{DR} . Even though some of the HVAC cannot provide a full shed during a long D_{DR} , their delayed schedule benefits highly from the decrease in the external temperatures.

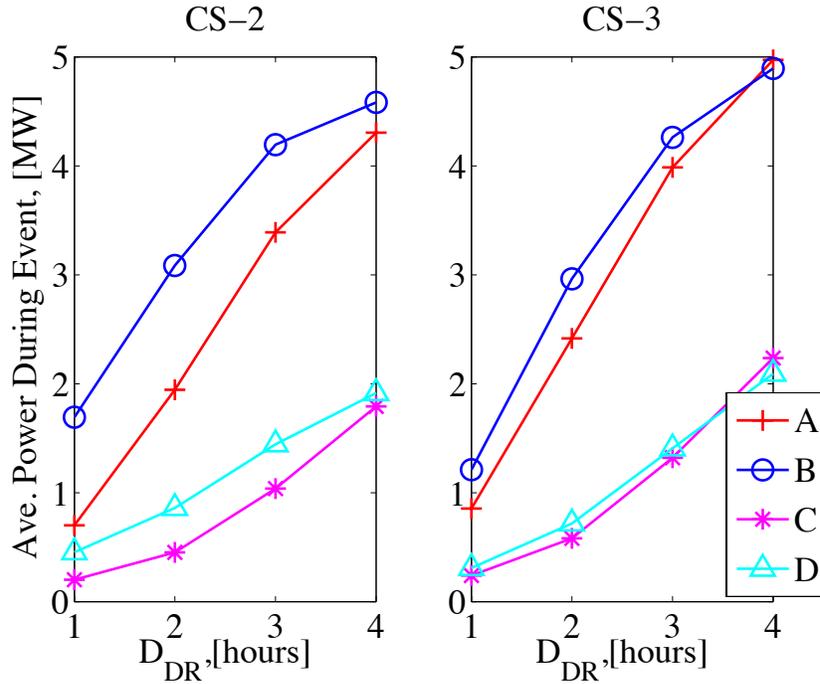


Figure 2.10: Average power during event for CS-2 and 3.

2.5 Chapter Conclusions

In this chapter, we have used a unique dataset to estimate and validate ARMA models that capture the thermal dynamics of individual HVAC units. Using these models, we investigated the potential of DR that can be provided with an aggregation of HVACs under certain conditions. We analyzed the variability in duration of the shed each HVAC unit can provide and we quantified the energy savings that can be achieved through various setpoint adjustment scenarios. Although, our results are obtained using data from residential units in the Texas area only, we believe that our study provides important insights into the aggregate behavior of HVAC units and their DR potential.

In the individual ARMA models, the estimated solar heat gains are often counter intuitive. They are relatively small and equally likely to be positive as they are neg-

ative, signifying the implausible interpretation that solar insolation acts as a cooling source in 50% of residences. We believe that this is an effect of using a solar-earth geometry model that does not account for cloud cover, and thus poorly represents variation in insolation inside the home. Instead, the regular diurnal pattern may correlate with occupancy, where residents are less likely to be home during the high insolation time (around noon). This may explain the negative coefficients for some homes, as high solar times correlate with low occupancy, reducing internal heat gains. Another reason may be the correlation between the outside air temperature and the regular diurnal insolation pattern. The colinearity of insolation with occupancy and/or outside air temperature is left as future work.

Assuming a constant coefficient of performance of the air conditioner also affects our estimates. We expect that this assumption causes us to over-predict the power required to achieve a certain level of cooling when the indoor/outdoor temperature differential is low and to under-predict the required power when the differential is high. Thus, on a very hot day, we should expect that the instantaneous savings from an DR event will be greater than predicted, as the HVAC unit will be operating at a lower COP.

Our validation suggests that we are over-predicting the thermal capacity of the building, causing our model to predict longer than observed durations of transient events following a setpoint change. This over-prediction could be caused by rounding errors in our low resolution data. For instance, when the room temperature is maintained at a setpoint, it is often measured to be constant when in fact we know it to be fluctuating within a dead band. Thus, the effect of HVAC cycling on internal temperature is masked by rounding during these times.

As discussed in Section 2.4, the proposed demand response potential estimation

strategy provided important insight into the response characteristics of an aggregate load population based on the data collected from individual households. We demonstrated how setpoint changes can be used to control a population of TCLs within an open-loop control strategy. However, as shown in Figure 2.7, open-loop control strategies often create a peak in power consumption following a DR event. Such a peak could negatively impact the power grid. In addition, individual modeling and control of TCLs might invite privacy concerns from end users. Furthermore, end users must be comfortable with the designated setpoint changes, and must have the ability to overwrite the setpoint changes if they feel uncomfortable. In the next chapter, we model an aggregation of TCLs to overcome some of these challenges.

Chapter 3

Modeling of Thermostatically Controlled Loads as Demand Response Agents

Even though the proposed ARMA models in Chapter 2 provide detailed insight into the behavior of individual TCLs under setpoint changes, modeling individual TCLs separately becomes intractable as the number of TCLs increases and individual models of TCLs are more likely to invite privacy concerns from the end users. Hence, in this chapter, we focus on the modeling of an aggregation of TCLs, such as heating, ventilation and air conditioning (HVAC) systems, refrigerators, and water heaters. We propose to make use of their thermal storage capacity under a centralized DLC strategy. These loads received significant attention because of their wide availability and because of their ability to be toggled ON/OFF without compromising end user satisfaction [13].

Different approaches to engage an aggregation of TCLs to the power grid via

This chapter is based on [31, 32, 33].

DLC have been investigated by many researchers. In particular, centralized control strategies have been widely studied (e.g., [31], [1], [35], [36], [38] and [39]), and this strategy is the main focus of this chapter. For example, the authors of [1] develop a linear time-invariant representation of a heterogeneous TCL population using a state-bin transition model, which we will use in this chapter as a benchmark. Similarly, in [38] a priority stack based algorithm to provide regulation services with an aggregate population of TCLs is presented. And in [40], the authors propose coupling household refrigerators with additional thermal storage to provide peak shaving. In our own previous work [31], we employ a Markov decision process (MDP) based approach to model the system dynamics of a centralized TCL control scenario using a state-bin transition model similar to [1].

In a centralized control scenario, the necessary communication between the TCLs and the central controller requires an additional sensing and information infrastructure, which raises privacy concerns and incurs additional costs [47, 1]. To address this issue, researchers have suggested using state estimation techniques that observe the aggregated power consumption of a population of TCLs and infer the current state of the population [48, 3, 2].

However, as addressed in [32], existing state estimation and control strategies are evaluated using simulations that are based on significant assumptions to model the operational use of TCLs. For example, thermal characteristics governing the interior temperature dynamics of individual TCLs and the ambient temperature are assumed to stay constant during the simulation. Furthermore, these parameters are often sampled from a normal and/or uniform distribution. Also, disturbances to the operation of the TCLs, such as door openings and changes of thermal mass in a TCL, are often ignored and/or included as Gaussian white noise in the interior

temperature evolution of TCLs [1, 3, 31, 2].

There has been limited work on exploring relaxations of these assumptions. In [49], different thermal resistance values for different load interaction scenarios are introduced. Specifically, the authors argue that user interaction plays an important role on thermal resistance (e.g., refrigerator door openings) and suggest using two different values during modeling. The alternation in between these two values is assumed to follow a homogeneous Poisson process. Yet, although these efforts suggest the need for an improvement in the modeling approach, further work is required to realistically model an aggregate TCL population in a way that considers time-dependent disturbances to their operation.

In this section, we relax many of the assumptions made so far and systematically explore the results of doing so. Our main motivation is to better understand whether the existing linear time-invariant (LTI) system formulations proposed in the literature (e.g., [1, 3]) are suitable for modeling the behavior of a collection of HRUs under realistic operating conditions. In particular, we focus on modeling the behavior of a collection of HRUs and evaluate the performance of these models under an open-loop control scenario with random actions. To model the behavior of HRUs under realistic operational conditions, building on our previous methodology proposed in [32], we fit probabilistic models to a year-long dataset of power measurements for 373 HRUs during regular operation. Specifically, in the current chapter, we use these models to create a realistic HRU population that is controlled under an open-loop strategy for the first time. We also expand the theoretical discussions in [2] by deriving the dynamic equations for the aggregate system using a bottom-up approach that results in a more flexible LTI system with a non-linear control rule. In particular, the proposed LTI system does not make any explicit assumptions about

the effects of the control actions on the state distributions. Finally, we quantify the plant-model mismatch and systematically evaluate the performance of both LTI systems using an open-loop control strategy.

In the next section, we summarize the existing individual and aggregate modeling strategies, specifically the benchmark system model from [3] (which we refer to as bLTI). Following that, we introduce our methodology. First, we describe the dataset, the individual load modeling strategy based on [32], and the aggregate load model of the HRUs. Then, we introduce the experiments that we conduct to evaluate the individual modeling strategy and the plant-model mismatch under different disturbance conditions. In Sections 3.3 and 3.4 we present and reflect on the results of these experiments.

3.1 Overview of Existing Models

In order to give the readers enough context to understand the remainder of the chapter, we provide a brief overview of the existing models used in the literature for both individual and aggregations of TCLs:

3.1.1 Modeling Individual TCLs

A two-state individual TCL model is commonly used in the literature to simulate these loads [26, 31, 1, 48, 3]. In such a model, the interior temperature of a TCL is modeled using heat transfer equations based on the thermal characteristics of the load and the TCL is assumed to have two states: *ON* and *OFF*. To create a heterogeneous population of TCLs, researchers sample thermal characteristics for each TCL from normal and/or uniform distributions and assume they are time-invariant over the course of the simulation. The difference equation governing the dynamics of each TCL can be generalized as follows, where $\theta_{i,t}$ is the internal temperature of

the i^{th} load at time t :

$$\theta_{i,t+1} = \alpha^{(i)}\theta_{i,t} + (1 - \alpha^{(i)})(\theta_i^a - m_{i,t}\theta_i^g) + \epsilon_{i,t} \quad (3.1)$$

The temperature gain, θ_i^g , is dependent on the *resistance* R_i and the *rated power* P_i^r of the appliance and is given as:

$$\theta_i^g = \begin{cases} R_i P_i^r, & \text{for cooling devices} \\ -R_i P_i^r, & \text{for heating devices} \end{cases} \quad (3.2)$$

In (3.1), $\epsilon_{i,t}$ is included as an error term on the internal temperature dynamics, and assumed to be white and Gaussian [3]. Also, the thermal characteristics of each TCL are expressed with the unitless parameter $\alpha^{(i)}$, defined as

$$\alpha^{(i)} = e^{-h/C_i R_i} \quad (3.3)$$

where C_i is the thermal capacitance and h is the simulation time step (i.e. the duration between t and $t + 1$).

A thermostat controls the internal temperature of each TCL around the set point θ_i^s by controlling the binary parameter m_i , which represents the current status of the appliance. The width of the thermostatic dead-band around θ_i^s is defined as δ_i for each TCL, such that $\theta_{i,t}$ for each TCL remains within the $[\theta^{low}\theta^{high}]$ range, where $\theta^{low} = \theta_i^s - \delta_i/2$ and $\theta^{high} = \theta_i^s + \delta_i/2$.

3.1.2 Modeling Aggregations of TCLs

Previous work used a state bin transition model [1] and [31] to capture the aggregate dynamics of a TCL population. Discrete temperature interval and status pairs rep-

represent individual TCL states, and the aggregated state probability mass distributed over these pairs represents the aggregate state of the TCL population, X_t . The elements of this vector (also called bins) can be defined as:

$$\begin{aligned} x_{k,t}^{ON} &= \mathbf{P}\{S_t = ON, I_t = k\} \\ x_{k,t}^{OFF} &= \mathbf{P}\{S_t = OFF, I_t = k\} \end{aligned} \quad (3.4)$$

where S_t is the current status of the HRU and I_t is the current temperature interval that the HRU is in at time t . The summation of all the elements $x_{k,t}^{ON}$ and $x_{k,t}^{OFF}$ at time t for all k must be equal to 1 for X_t to be a valid probability mass distribution. The vector representation of the state is then:

$$X_t = [x_{1,t}^{OFF}, \dots, x_{N,t}^{OFF}, x_{1,t}^{ON}, \dots, x_{N,t}^{ON}] \quad (3.5)$$

Due to individual TCL dynamics (i.e. heat exchange due to difference between $\theta_{a,t}$ and $\theta_{i,t}$ and the thermostatic control actions), each TCL moves through these bins as depicted in Figure 3.1. Therefore the probability mass distributed between the discrete bins forming X_t is time dependent. In Figure 3.1, θ^{low} represents the low end of the dead-band and θ^{high} represents the high end of the dead-band.

The TCL switching strategy in [1], [48], [3], [31] and [2] uses a control signal that includes the switching probabilities that are sent to HRUs in discrete temperature and status pairs depicted in Figure 3.1. The switching probability D_t is composed of individual switching probabilities per temperature interval I_t and per switching direction as follows:

$$\begin{aligned} d_{k,t}^{ON|OFF} &= \mathbf{P}\{S_{t+1} = ON | S_t = OFF, I_t = k, cont.\} \\ d_{k,t}^{OFF|ON} &= \mathbf{P}\{S_{t+1} = OFF | S_t = ON, I_t = k, cont.\} \end{aligned} \quad (3.6)$$

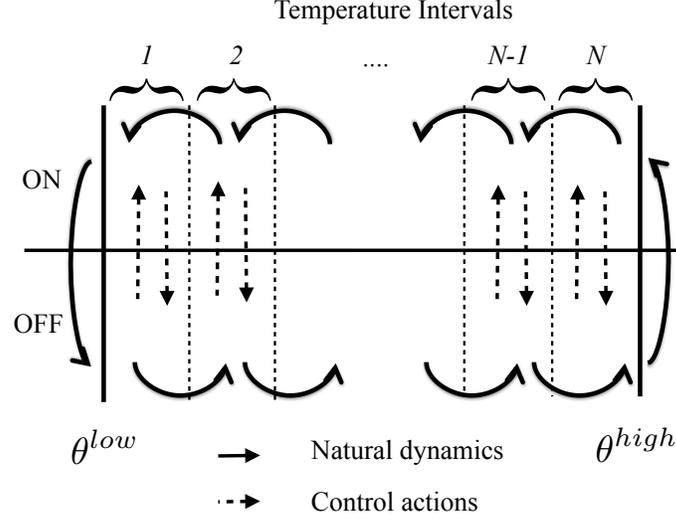


Figure 3.1: State bin transition model adapted from [1] and [2]

The vector representation of switching probabilities is thus:

$$D_t = [d_{1,t}^{ON/OFF}, \dots, d_{N^{int},t}^{ON/OFF}, d_{1,t}^{OFF/ON}, \dots, d_{N^{int},t}^{OFF/ON}] \quad (3.7)$$

The benchmark LTI (bLTI) system is adopted from [48] and [3] and is obtained using the Markov chain model introduced in [1]. The natural dynamics of the system are captured using a transition matrix (\mathbf{A}^{bLTI}), whose elements denote the transition probability between bins. Previous work defines the control signal as the mass exchange between intervals calculated by $X_t \circ D_t$ (where \circ is the Hadamard product operator). Here we explicitly include this term in the equation and use the switching probabilities D_t as the control actions for clarity.

Also, it is assumed that the \mathbf{B}^{bLTI} has the following structure to ensure that the next state vector is a valid probability mass distribution and that the probability

mass exchange occurs within a temperature interval:

$$\mathbf{B}^{\text{bLTI}} = \begin{bmatrix} -\mathbb{1}_{N^{int}} & \mathbb{1}_{N^{int}} \\ \mathbb{1}_{N^{int}} & -\mathbb{1}_{N^{int}} \end{bmatrix} \quad (3.8)$$

where $\mathbb{1}_{N^{int}}$ is the identity matrix of size N^{int} .

In [48] and [3] the authors ensure that each temperature interval receives a single switching probability since bi-directional actions do not provide any benefits in terms of realizability of the desired change in the state and they increase the overall number of switchings. Finally, the adopted linear system can be expressed using X_t as defined in (3.5):

$$\begin{aligned} X_{t+1} &= \mathbf{A}^{\text{bLTI}} X_t + \mathbf{B}^{\text{bLTI}} X_t \circ D_t \\ Y_t &= \mathbf{C} X_t \end{aligned} \quad (3.9)$$

Generally, A^{bLTI} is estimated using a simulated population of TCLs with pre-determined characteristics and keeping track of switchings between different status and interval pairs in time. For notational purposes, we include the measurement equation $Y_t = \mathbf{C} X_t$ in (3.9) to have a complete state space representation. However, in this work we mainly focus on state prediction and plant-model mismatch.

3.2 Methodology

Now we discuss the relaxations we have made to the assumptions made by the models presented above. To maintain consistency with the way we have introduced the concepts so far, we will continue to separate the discussion of the individual TCL models and the aggregation of TCLs. In other words, we will introduce our approach, experiments and results for these two types of models in separate sections

Table 3.1: Dataset used for HRUs [4]

<i>Sampling Period</i>	<i>Collection Duration</i>	<i>Refrigeration Unit</i>	<i>Total Number</i>
2 Minutes	27 Days	Freezer Upright	89
	28 Days	Chest Freezer	39
	28 Days	Refrigerator	101
	27 Days	Fridge + Freezer	39
	28 Days	Beer/Wine Chiller	3

throughout the remainder of the manuscript.

3.2.1 Dataset Description

The Household Energy Survey of the UK includes data on electrical power demand of 251 households over the period from May 2010 to July 2011. Specifically, it consists of multiple datasets with different numbers of households, collection durations, and sampling rates [4].

In this study, we use the energy demand measurements collected every 2 minutes from refrigeration units over one month within the May 2010 to July 2011 period. We use data taken from 373 different refrigeration units present in 214 households. Table 3.1 shows the characteristics of this dataset along with the types and quantities of refrigeration units. The energy demand is measured in deciwatt hours, and to estimate the rated power of each refrigeration unit, we assumed the rated power of a unit remains constant between samples.

3.2.2 Modeling Individual TCLs

Instead of sampling the thermal characteristics in the way previous studies have done, we use the modeling strategy we proposed in [32]. Specifically, we make use of a hierarchical model of T_i^{ON} and T_i^{OFF} which are defined as the time that it takes for each HRU to go from the upper margin of the thermostatic dead-band to the lower

margin and vice versa, respectively. We assume that the random variables T_i^{ON} and T_i^{OFF} follow a 3 parameter Weibull distribution for each load i (a comparison of the goodness of fit of other distributions can be found in [32]). The probability density function for a 3 parameter Weibull distribution is given as follows:

$$f(t) = \begin{cases} \frac{\beta}{\eta} \left(\frac{t-\gamma}{\eta}\right)^{\beta-1} e^{-\left(\frac{t}{\eta}\right)^\beta}, & t \geq 0 \\ 0, & t < 0 \end{cases} \quad (3.10)$$

where η is the scale parameter, β is the shape parameter and γ is the location parameter.

In contrast to our previous work [32], where we used a two-parameter Weibull distribution, we introduce the location parameter, γ . The main reason behind introducing the location parameter is that it allows including a lower bound for the values of T_i^{ON} and T_i^{OFF} . The γ parameter shifts the Weibull distribution along the t-axis by γ and it is taken as the data sampling rate from the HES dataset (i.e., 2 minutes) [4]. This means that the probability distributions are defined for values of T_i^{ON} and T_i^{OFF} higher than or equal to the sampling rate of the HES dataset. Estimating the Weibull parameters $(\beta_i^{OFF}, \eta_i^{OFF})$ and $(\beta_i^{ON}, \eta_i^{ON})$ that represents the distributions of T_i^{ON} and T_i^{OFF} helps us recreate the monitored HRUs in simulation. In order to generalize the modeling strategy further, we look at the distributions of parameters $(\beta_i^{OFF}, \eta_i^{OFF})$ and $(\beta_i^{ON}, \eta_i^{ON})$ among all HRUs. After evaluating different families of distributions based on the Akaike and Bayesian information criteria (AIC and BIC, respectively), we arrived at the distributions given in Table 3.2.

The individual modeling strategy of HRUs can now be described as follows. We first sample the hyper-parameters of the distributions that govern the varying T_i^{ON} and T_i^{OFF} values for different operational conditions using the distributions given in

Table 3.2: Hyper-Parameters of T_i^{ON} and T_i^{OFF}

<i>Parameter</i>	<i>Distribution</i>	<i>Hyper-Parameters</i>	<i>Values</i>
β_i^{ON}	Generalized Pareto	k : tail index	$k = 1.879$
		σ : scale	$\sigma = 2.232$
		θ : location	$\theta = 0.793$
η_i^{ON}	Log-normal	μ : mean	$\mu = 2.885$
		σ : scale	$\sigma = 0.881$
β_i^{OFF}	Generalized Pareto	k : tail index	$k = 0.433$
		σ : scale	$\sigma = 5.529$
		θ : location	$\theta = 0.540$
η_i^{OFF}	Gamma	a : shape	$\mu = 2.231$
		b : scale	$b = 20.704$

Table 3.2. Then, we store these parameters for each HRU i that we are simulating. Lastly, we use these hyper-parameters to sample different T_i^{ON} and T_i^{OFF} values from a Weibull distribution for each HRU i . Re-sampling different T_i^{ON} and T_i^{OFF} values for each refrigerator allows for varying operational conditions, and allows the disturbances to individual HRUs to be included in the simulation. Whenever a different value for T_i^{ON} and T_i^{OFF} is sampled, the corresponding thermal characteristics are obtained using (3.13) and (3.14) adopted from [32]. We assume a constant ambient temperature since the changes in the ambient temperature are reflected in the hierarchical model creation.

Using (3.1) we can write the T_i^{ON} and T_i^{OFF} in terms of the temperature set point $\theta_{set,i}$, the thermostatic dead-band width δ_i , the ambient temperature θ_i^a , and the rated power P_i^r .

$$T_i^{ON} = -R_i C_i \ln\left(\frac{\theta_i^s - \delta_i/2 - \theta_i^a + P_i^r R_i}{\theta_i^s + \delta_i/2 - \theta_i^a + P_i^r R_i}\right) \quad (3.11)$$

$$T_i^{OFF} = -R_i C_i \ln\left(\frac{\theta_i^s + \delta_i/2 - \theta_i^a}{\theta_i^s - \delta_i/2 - \theta_i^a}\right) \quad (3.12)$$

Now, using (3.11) and (3.12), thermal resistance R_i for each HRU is

$$R_i = \frac{(\theta_i^a - \theta_i^s)(K_i - 1) - \delta_i/2(K_i + 1)}{P_i^r(K_i - 1)} \quad (3.13)$$

where K_i is defined as

$$K_i = \left(\frac{\theta_i^s + \delta_i/2 - \theta_i^a}{\theta_i^s - \delta_i/2 - \theta_i^a}\right)^{\frac{T^{ON}}{T_i^{OFF}}} \quad (3.14)$$

The following similar expression is derived for the thermal capacitance C_i .

Experiments

To evaluate how realistic the variability of the R_i and C_i values assumed in the literature is, we refer to values reported by Mathieu [5] and create appliance populations that consist of a hundred thousand appliances using both modeling approaches. For the first population (P1), we sample the Weibull parameters β_i^{OFF} , β_i^{ON} , η_i^{OFF} and η_i^{ON} using the hyper-parameters obtained from the Table 3.2. Using these variables we sample a T^{ON} and a T_i^{OFF} value from the corresponding Weibull distributions for each unit. Then we calculate the R_i and C_i values using (3.13) and (3.14). In order to be able to compare the variability in R_i and C_i values we assume that P_i^r , $\theta_{(i,a)}$, δ_i and $\theta_{i,set}$ distributions are as given in Table 3.3. The resulting parameters and their characteristics are described in Table 3.4. For the second population (P2), we randomly sample R_i and C_i parameters using the distributions given in Table 3.3. Finally, we compare the sample means \hat{R} and \hat{C} , the standard deviations σ^R and σ^C , and the coefficients of variation $c_v^R = \sigma^R/\hat{R}$ and $c_v^C = \sigma^C/\hat{C}$ between the two populations.

Table 3.3: HRU parameters obtained from [5]

	R_i ($^{\circ}\text{C}$)	C_i ($\text{kWh}/^{\circ}\text{C}$)	P_i^r (kW)	θ_i^a ($^{\circ}\text{C}$)	δ_i ($^{\circ}\text{C}$)	$\theta_{i,set}$ ($^{\circ}\text{C}$)
Distribution	Uniform	Uniform	Uniform	Constant	Uniform	Uniform
Range	80-100	0.4-0.8	0.2-1.0	20	1-2	1.7-3.3

Table 3.4: Individual simulation parameters for HRUs i

<i>Parameter</i>	<i>Distribution Characteristics</i>	<i>Range</i>
θ_i^a	Constant	20°C
T_i^{OFF}	Weibull	obtained via Table 3.2
T_i^{ON}	Weibull	obtained via Table 3.2
h	Constant	4 seconds
P_i^r	Uniform	$[0.1, 0.2] \text{ kW}$
δ	Constant	2°C
θ_i^s	Constant	2°C
N^{int}	Constant	3
R	Calculated from T_i^{ON} and T_i^{OFF}	
C	Calculated from T_i^{ON} and T_i^{OFF}	

3.2.3 Modeling Aggregations of TCLs

Using the definitions of X_t and D_t given in (3.5) and (3.7), we present a state transition function $\mathfrak{T}(X_t, D_t)$ as part of a Markov Decision Process framework in [31] and [2]. In this chapter, we use this state transition function ($\mathfrak{T}(X_t, D_t)$) to systematically derive a time-invariant model which is state-dependent in control actions.

First, in order to separate between the thermostats in each HRU and the main controller, we define λ_k as the probability of switching due to thermostatic actions at each temperature interval in both switching directions:

$$\begin{aligned}
 \lambda_k^{ON/OFF} &= P\{S_{t+1} = ON | S_t = OFF, I_t = k, \text{natural}\} \\
 \lambda_k^{OFF/ON} &= P\{S_{t+1} = OFF | S_t = ON, I_t = k, \text{natural}\}
 \end{aligned} \tag{3.15}$$

Since the thermostatic actions are expected to yield a change in the current status of the appliance, they are not expected to change the HRU's temperature interval. Although this is also true for control actions, a new set of parameters is defined to differentiate between the probabilities to switch conditioned on the natural dynamics ($A_{k,i}$) or the control actions ($O_{k,i}$). These terms are defined for each switching direction, each current temperature interval $I_t = k$ and for every possible next temperature interval $I_{t+1} = i$ as:

$$\begin{aligned} O_{k,i}^{ON/OFF} &= P\{I_{t+1} = i | S_{t+1} = ON, S_t = OFF, I_t = k, controller\} \\ A_{k,i}^{ON/OFF} &= P\{I_{t+1} = i | S_{t+1} = ON, S_t = OFF, I_t = k, natural\} \end{aligned} \quad (3.16)$$

Similar definitions can also be given for cases with OFF to ON switching. For $A_{k,i}$ terms, we also define $A_{k,i}^{OFF/OFF}$ and $A_{k,i}^{ON/ON}$ to capture the natural dynamics for the HRUs that are not controlled during the simulations.

Now, the system behavior can be described as follows. The controller sends out a decision vector D_t of switching probabilities $d_{k,t}$. If the appliance decides to switch based on the probability it received, the transition to another temperature interval is determined by the corresponding $O_{k,i}$ terms; otherwise, it is determined by the $A_{k,i}$ terms. The HRUs in these bins are assumed to know their current temperature and statuses so that they switch based on the probability received in the control signal.

Using these definitions, the state transition function $\mathfrak{T}(X_t, D_t)$ can be shown to

have the following form:

$$\begin{aligned}
x_{i,t+1}^{ON} = & \sum_{k=1}^N \left[A_{k,i}^{ON/OFF} \lambda_k^{ON/OFF} (1 - d_{k,t}^{ON/OFF}) x_{k,t}^{OFF} \right. \\
& + O_{k,i}^{ON/OFF} d_{k,t}^{ON/OFF} x_{k,t}^{OFF} \\
& \left. + A_{k,i}^{ON/ON} (1 - \lambda_k^{OFF/ON}) (1 - d_{k,t}^{OFF/ON}) x_{k,t}^{ON} \right]
\end{aligned} \tag{3.17}$$

A similar equation can be formed to obtain $x_{i,t+1}^{OFF}$ for all i .

$$\begin{aligned}
x_{i,t+1}^{OFF} = & \sum_{k=1}^N \left[A_{k,i}^{OFF/ON} \lambda_k^{OFF/ON} (1 - d_{k,t}^{OFF/ON}) x_{k,t}^{ON} \right. \\
& + d_{k,t}^{OFF/ON} O_{k,i}^{OFF/ON} x_{k,t}^{ON} \\
& \left. + A_{k,i}^{OFF/OFF} (1 - \lambda_k^{ON/OFF}) (1 - d_{k,t}^{ON/OFF}) x_{k,t}^{OFF} \right]
\end{aligned} \tag{3.18}$$

The state space representation of the above equation is as follows:

$$X_{t+1} = \mathbf{A}^{\text{LTI}} X_t + \mathbf{B}^{\text{LTI}} X_t \circ D_t \tag{3.19}$$

where X_t and D_t are as defined in (3.5) and (3.7), and \mathbf{A}^{LTI} and \mathbf{B}^{LTI} are $2N^{\text{int}} \times 2N^{\text{int}}$ matrices shown on the following page (3.20) and (3.21), respectively.

As opposed to the \mathbf{B}^{bLTI} adopted from [3] and defined in (3.8), the system representation given in (3.19) is not based on a pre-determined behavior of the system under control actions, as can be seen by the presence of the O terms in the \mathbf{B}^{LTI} matrix shown in (3.21). Hence, it adds another degree of freedom to the model.

$$\mathbf{A}^{\text{LTI}} = \begin{bmatrix} (1 - \lambda_1^{\text{ON/OFF}})A_{1,1}^{\text{OFF/OFF}} & \dots & (1 - \lambda_1^{\text{ON/OFF}})A_{N^{int},1}^{\text{OFF/OFF}} & \lambda_1^{\text{OFF/ON}}A_{1,1}^{\text{OFF/ON}} & \dots & \lambda_{N^{int}}^{\text{OFF/ON}}A_{N^{int},1}^{\text{OFF/ON}} \\ \vdots & & \vdots & \vdots & & \vdots \\ (1 - \lambda_1^{\text{ON/OFF}})A_{1,N^{int}}^{\text{OFF/OFF}} & \dots & (1 - \lambda_1^{\text{ON/OFF}})A_{N^{int},N^{int}}^{\text{OFF/OFF}} & \lambda_1^{\text{OFF/ON}}A_{1,N^{int}}^{\text{OFF/ON}} & \dots & \lambda_{N^{int}}^{\text{OFF/ON}}A_{N^{int},N^{int}}^{\text{OFF/ON}} \\ \lambda_1^{\text{ON/OFF}}A_{1,1}^{\text{ON/OFF}} & \dots & \lambda_{N^{int}}^{\text{ON/OFF}}A_{N^{int},1}^{\text{ON/OFF}} & (1 - \lambda_1^{\text{OFF/ON}})A_{1,1}^{\text{ON/ON}} & \dots & (1 - \lambda_{N^{int}}^{\text{OFF/ON}})A_{N^{int},1}^{\text{ON/ON}} \\ \vdots & & \vdots & \vdots & & \vdots \\ \lambda_1^{\text{ON/OFF}}A_{1,N^{int}}^{\text{ON/OFF}} & \dots & \lambda_{N^{int}}^{\text{ON/OFF}}A_{N^{int},N^{int}}^{\text{ON/OFF}} & (1 - \lambda_1^{\text{OFF/ON}})A_{1,1}^{\text{ON/ON}} & \dots & (1 - \lambda_{N^{int}}^{\text{OFF/ON}})A_{N^{int},N^{int}}^{\text{ON/ON}} \end{bmatrix} \quad (3.20)$$

$$\mathbf{B}^{\text{LTI}} = \begin{bmatrix} A_{1,1}^{\text{OFF/OFF}}(\lambda_1^{\text{ON/OFF}} - 1) & \dots & A_{1,N^{int}}^{\text{OFF/OFF}}(\lambda_1^{\text{ON/OFF}} - 1) & O_{1,1}^{\text{OFF/ON}} - \lambda_1^{\text{OFF/ON}}A_{1,1}^{\text{OFF/ON}} & \dots & O_{1,N^{int}}^{\text{OFF/ON}} - \lambda_1^{\text{OFF/ON}}A_{1,N^{int}}^{\text{OFF/ON}} \\ \vdots & & \vdots & \vdots & & \vdots \\ A_{N^{int},1}^{\text{OFF/OFF}}(\lambda_1^{\text{ON/OFF}} - 1) & \dots & A_{N^{int},N^{int}}^{\text{OFF/OFF}}(\lambda_{N^{int}}^{\text{ON/OFF}} - 1) & O_{N^{int},1}^{\text{OFF/ON}} - \lambda_{N^{int}}^{\text{OFF/ON}}A_{N^{int},1}^{\text{OFF/ON}} & \dots & O_{N^{int},N^{int}}^{\text{OFF/ON}} - \lambda_{N^{int}}^{\text{OFF/ON}}A_{N^{int},N^{int}}^{\text{OFF/ON}} \\ O_{1,1}^{\text{ON/OFF}} - \lambda_1^{\text{ON/OFF}}A_{1,1}^{\text{ON/OFF}} & \dots & O_{1,N^{int}}^{\text{ON/OFF}} - \lambda_1^{\text{ON/OFF}}A_{1,N^{int}}^{\text{ON/OFF}} & A_{1,1}^{\text{ON/ON}}(\lambda_1^{\text{OFF/ON}} - 1) & \dots & A_{1,N^{int}}^{\text{ON/ON}}(\lambda_1^{\text{OFF/ON}} - 1) \\ \vdots & & \vdots & \vdots & & \vdots \\ O_{N^{int},1}^{\text{ON/OFF}} - \lambda_{N^{int}}^{\text{ON/OFF}}A_{N^{int},1}^{\text{ON/OFF}} & \dots & O_{N^{int},N^{int}}^{\text{ON/OFF}} - \lambda_{N^{int}}^{\text{ON/OFF}}A_{N^{int},N^{int}}^{\text{ON/OFF}} & A_{N^{int},1}^{\text{ON/ON}}(\lambda_{N^{int}}^{\text{OFF/ON}} - 1) & \dots & A_{N^{int},N^{int}}^{\text{ON/ON}}(\lambda_{N^{int}}^{\text{OFF/ON}} - 1) \end{bmatrix} \quad (3.21)$$

System Identification

To train the system matrices A and O and the system vector λ that define the state space matrices \mathbf{A}^{LTI} and \mathbf{B}^{LTI} , we simulate a population of HRUs as described in Section 3.2 for a total duration of 24 hours. After the simulation has run for 12 hours, we send a random action D_t to the simulated plant at each time step t . The reasons for this are two-fold. First, we would like to capture the behavior of the simulated plant under control actions in order to represent it by the O matrices. Second, we would like to ensure that the population has reached a steady state operation. During the course of the simulation, we tally each transition from an interval I_t to status S_t pair and the causing action (i.e., due to centralized controller or the local thermostat). Finally, we estimate the probabilities associated with each transition.

Experiments

To evaluate the performance of the proposed LTI system and the bLTI system adopted from [3], we create a population of 1000 HRUs simulated with the individual modeling strategy described in Section 3.2.2. We create an open-loop control scenario, where we send random control actions to a simulated HRU population. For every action sent, we use the proposed LTI and the bLTI models to predict the state in the next time step. That is, at each time step, we assume both models have perfect information on the control action and the current state of the HRU population. We then predict the next state using both models. A schematic describing the experiment is shown in Figure 3.2.

To better understand the effect of disturbances to individual units on the plant-model mismatch, we control how frequently resampling of the thermal parameters

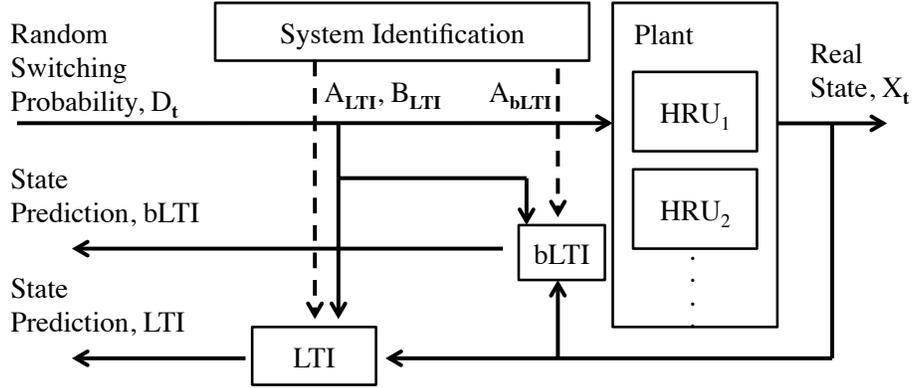


Figure 3.2: Plant-model mismatch evaluation

occurs in each HRU by introducing the term $p_{d,i}$, the *density of disturbance*, as the probability of a disturbance to the HRU i . A disturbance implies resampling the T_i^{ON} and T_i^{OFF} values using the saved hyper-parameters for HRU i . In addition to $p_{d,i}$, we also control how frequently each appliance is switched by the centralized controller by introducing a forcing parameter, f , which is defined as the maximum switching probability that any temperature interval receives from the main controller.

To evaluate the effects of the *density of disturbance*, $p_{d,i}$, forcing parameter f and the simulation time step h on the plant-model mismatch –how well each model captures the behavior of the plant– we ran three different scenarios, S1-S3, with varying $p_{d,i}$ values corresponding to the expected duty cycle duration, $T_i^{ON} + T_i^{OFF}$, (i.e., once every hour), above the expected duty cycle duration (i.e., once every half hour) and below the expected duty cycle duration (i.e., once every two hours), respectively. In each scenario, different values for the forcing parameter, f and the time step h are also investigated. The characteristics of different parameters for each scenario are given in more detail in Table 3.5.

Table 3.5: Different scenarios with varying $p_{d,i}$, h and f

<i>Scenario</i>	<i>Density of Disturbance</i>	<i>Forcing Parameter</i>	<i>time step</i>
S1	$p_{d,i} = 20/900$ (every 30 minutes)	$f = 1$	4, 15, 60 seconds
		$f = 0.75$	
		$f = 0.5$	
		$f = 0.25$	
S2	$p_{d,i} = 10/900$ (every 60 minutes)	$f = 1$	4, 15, 60 seconds
		$f = 0.75$	
		$f = 0.5$	
		$f = 0.25$	
S3	$p_{d,i} = 4/900$ (every 120 minutes)	$f = 1$	4, 15, 60 seconds
		$f = 0.75$	
		$f = 0.5$	
		$f = 0.25$	

Table 3.6: Estimated parameters for HRU populations with a hundred thousand units: P1 and P2

Population	\hat{R}_i ($^{\circ}\text{C}/\text{kW}$)	\hat{C}_i ($\text{kWh}/^{\circ}\text{C}$)	σ^R ($^{\circ}\text{C}/\text{kW}$)	σ_C ($\text{kWh}/^{\circ}\text{C}$)	c_v^R	c_v^C
P1	419.41	0.07	9205.7	0.07	21	1
P2	90	0.6	175.76	0.12	0.06	0.19

3.3 Results

3.3.1 Individual TCL Behavior

Although the estimated \hat{R} and \hat{C} values for different populations provide little information when compared, the coefficient of variations given in Table 3.6 show that the assumptions made for the parameter values and distributions in P2 result in a squeezed and less varying distribution. Appliances simulated based on assumptions followed by P2 will provide similar availability for demand response services in each cycle. Further, the dispersion around the estimated sample mean, measured by the c_v^R and c_v^C , appears to be more significant for P1, indicating a relatively stretched distribution. We believe this may be due to the time dependent disturbances incor-

porated in the proposed model and the variety of different refrigeration unit types in the dataset.

3.3.2 Plant-Model Mismatch

Figure 3.3 shows the root mean squared error (RMSE) between the predicted states and the actual states X_t . The errorbars mark the $\pm\sigma_{RMSE}$ values (i.e., values that are above and below by a standard deviation). It is possible to observe that: for both of the systems, as the forcing parameter f increases, the RMSE increases independent of the time step h and the disturbance density $p_{d,i}$. It is also possible to observe that within the same h values, there is no significant change in RMSE value as the $p_{d,i}$ increases.

The results also suggest that a faster sampling of the state variables may mitigate the effects of invalid assumptions about the disturbances and time-invariant properties of the individual HRUs that the bLTI system is making. However, when this faster sampling is not possible, the proposed LTI system reduces the plant-model mismatch.

In Figure 3.4, the \mathbf{B}^{bLTI} matrices obtained from S1 where $h = 4 \text{ seconds}$ and S3 where $h = 60 \text{ seconds}$ are shown as gray-scaled images using the color-mapping depicted. To understand the differences between the \mathbf{B}^{LTI} and the \mathbf{B}^{bLTI} with varying $p_{d,i}$, f and h , we have also included a reference \mathbf{B}^{bLTI} which is identical in any cases. It is possible to observe that the dissimilarity between the \mathbf{B}^{bLTI} and the \mathbf{B}^{LTI} becomes visible at the marginal intervals when h and f are higher, which suggests that the additional degree of freedom provided by the training step, helps decreasing the RMS error.

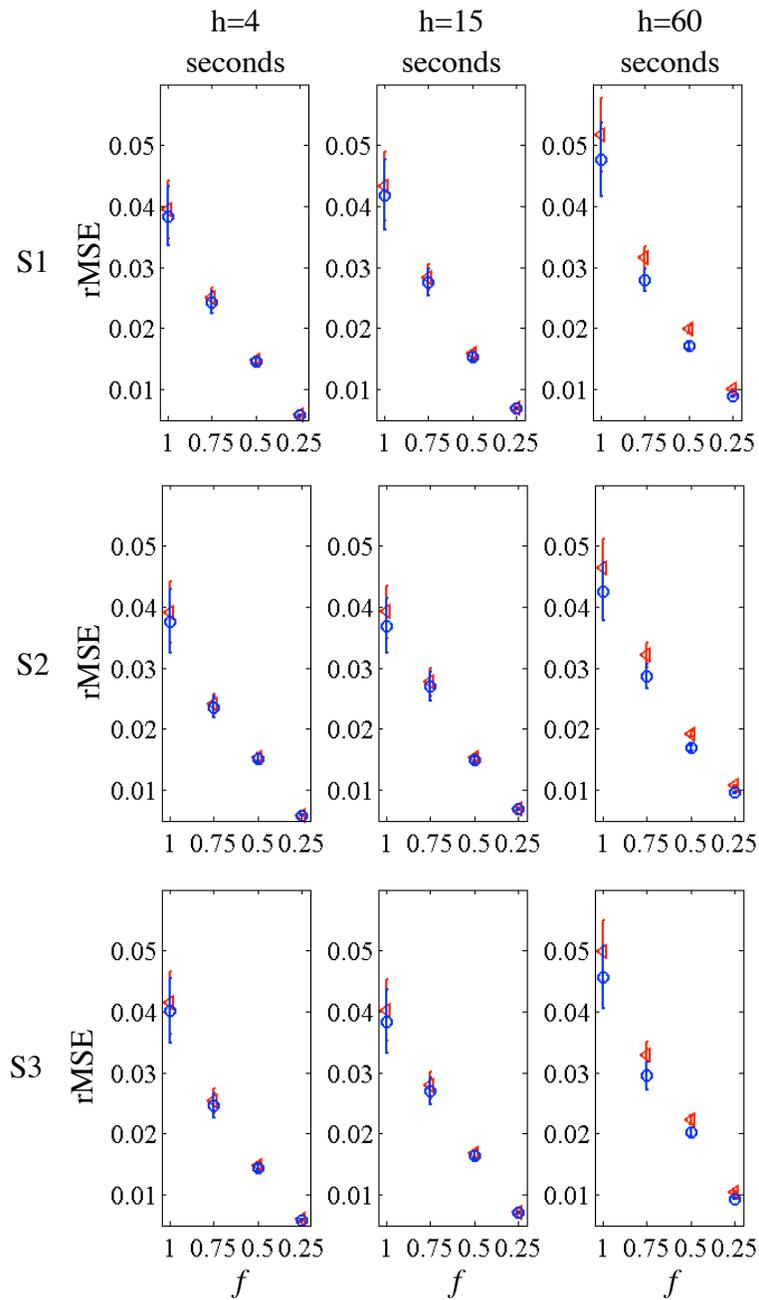


Figure 3.3: Plant model mismatch using the proposed LTI strategy and the bLTI model adopted from [3] where f is the forcing parameter and each row corresponds to result of scenarios S1,S2 and S3

3.4 Chapter Conclusions

The main motivation of this chapter is to better understand if the existing linear time-invariant (LTI) system formulations proposed in the literature (e.g., [1, 3]) are

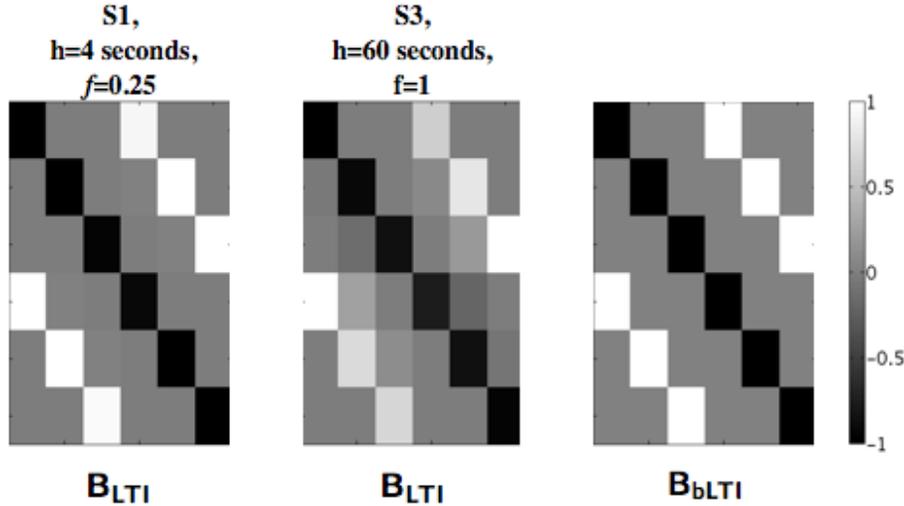


Figure 3.4: Training results for \mathbf{B}^{LTI} and \mathbf{B}^{bLTI} . The vertical color bar displays the mapping between the value of the each element in the matrices and the corresponding color used to represent it.

capable of modeling the behavior of a collection of HRUs under realistic operating conditions. Specifically, we investigate the assumptions on time-invariant thermal parameters and the strict definition of the \mathbf{B}^{bLTI} in the benchmark strategy that enforces a mass exchange within the temperature interval.

Our results suggest that, if the plant is simulated close to real-time, these assumptions are valid. However, neither information on R and C values of the loads nor historical energy consumption data from these loads may be available. In addition, there may be limits on the communication infrastructure during training, hence the state observations of the population may be unavailable close to real-time. The proposed LTI strategy brings extra value in such cases.

Even though the timing requirements of different ISOs/RTOs participating in ancillary services differ between 2-10 seconds [19, 50, 51], the requirement decreases to 5 to 10 minutes for services such as load following [19]. Depending on the service to be provided (i.e. the timing requirement) and the modeling accuracy needed, the

proposed LTI system can be used to better capture the dynamics of the aggregate HRU population.

Finally, the proposed LTI strategy, where \mathbf{B}^{LTI} is trained by taking random actions on the population of HRUs, has no computational burden over the benchmark LTI system. Our results suggest that the benchmark LTI system acts as a boundary case for the proposed LTI system and the additional degree of freedom brought by \mathbf{B}^{bLTI} brings extra value for specific cases when the simulation time steps are higher (e.g., load following) and non-linearities become more evident.

Chapter 4

State Estimation of Thermostatically Controlled Loads for Demand Response

In Chapter 3, we focused on an aggregate modeling strategy for a heterogeneous TCL population. In this chapter, we shift our attention to the upstream communication between the TCLs and the main controller, and the characteristics of the underlying cyberinfrastructure. In particular, we first introduce a moving horizon state estimator (MHSE) with constraints to infer the state of a heterogeneous TCL population, assuming the aggregate power consumption of TCLs is available to the main controller. We compare this estimator to a Kalman filter obtained from the literature, which assumes a linear model without constraints. The results show that some improvement is possible in scenarios where loads are expected to be toggled frequently.

Then, we relax our assumption on the availability of the aggregate power mea-

This chapter is partially based on [31, 2, 33].

measurements of TCLs, and assume that only a neighborhood-level aggregate power measurement is available to the main controller. We adopt our MHSE strategy to infer the states of TCLs when other loads are present in the aggregate measurements obtained at a neighborhood level. We use electrical power demand data collected from 251 households [4] to create a virtual neighborhood and test our state estimation strategy using the aggregate power consumption of all households measured from a neighborhood-level aggregation point.

4.1 Introduction

In recent years, there has been an increased interest in using TCLs, such as HVAC systems, refrigerators, and electric water heaters, to leverage demand-side mechanisms for improving grid-level efficiency, stability, and other properties of the electric power system. The availability of TCLs in most households and their ability to be toggled ON/OFF with minimal hindrance to their end use functions make them favorable for a control mechanism capable of providing a quicker response to intermittencies caused by renewable energy sources [13].

To realize the potential of using a population of TCLs to support grid-level objectives, the centralized control strategies that are proposed typically make assumptions that raise privacy concerns and require additional sensing and communication infrastructure. These concerns and additional sensing requirements are due to the necessary information exchange regarding thermal characteristics and state information between the loads and the main controller [47, 1]. To address these concerns, researchers have proposed different state estimation techniques to obtain the state information of a population of loads based on observations of aggregate power consumption [48, 3].

The contributions of this chapter are as follows: First, we propose a moving horizon mean squared error (MSE) state estimator with constraints. We refer to this estimator as the *moving horizon state estimator* (MHSE) based on the system dynamics developed in [31, 33], as described in section 3.1.2. For the MHSE we assume that aggregate power measurements of a heterogeneous TCL population are available to the main controller, and the estimator uses this information to infer the underlying population states.

The proposed MHSE provides an improved state estimation performance under conditions where the TCLs are expected to be toggled more frequently in comparison to the Kalman filtering based approach obtained from [3]. In particular, more frequent toggling is more likely when there are abrupt changes in the reference power that is to be followed, and when the system is expected to follow a reference power consumption over a certain period of time (e.g., load following).

Next, we leverage the contextually supervised source separation technique proposed in [52] to develop a state estimation strategy, where power measurements obtained at a neighborhood level distribution feeder are used to infer the state of the TCL population. We refer to this estimator as the *contextually supervised state estimator* (CSSE). As a proof-of-concept demonstration, we compare the performance of the CSSE to a simple prediction strategy with random initial starting points. In cases where the CSSE has perfect information on the thermal characteristics of the TCLs in the aggregate population, our results show that the CSSE outperforms the simple prediction strategy. Our results also suggest that as the proportion of non-TCL loads to TCL loads in the neighborhood level measurement point grows, the performance of the CSSE decays. In addition, we investigate a case where the thermal properties of the TCLs in the load population are not available

to the CSSE. Instead, we use the thermal properties proposed in the literature to perform system identification. When such a system model is used, the performance of the CSSE decays significantly, which suggests that the performance of the CSSE highly depends on the availability of a reliable model.

This chapter is organized as follows. In Section 4.2, we introduce the related work. In Section 4.3, we explain the individual and population modeling approaches for TCLs and the control signal used to realize actions in the heterogeneous TCL population. In Section 4.3.1, the MHSE building on [31] and the Kalman filtering framework, based on [3], is described and system identification techniques are explained. In Section 4.4, the resulting filters are compared using observations from identical plants. Each scenario is based on a heterogeneous TCL population simulated using unique appliances whose *capacitance, resistance and rated power consumption* are randomly selected from a uniform distribution within a given range. In Section 4.5, we describe the contextually supervised state estimator. In Section 4.7, we introduce the dataset used in the second part of this chapter and the load simulation strategy necessary to create a test bed for CSSE. In Section 4.8, we discuss the assumptions made in the CSSE and introduces two different case studies investigating the performance of CSSE with varying input parameters. Finally, conclusions and discussions are given in Section 4.9.

4.2 Related Work

A number of researchers have investigated methodologies to engage TCLs to the power grid via direct load control. Centralized control strategies using aggregated load populations are studied in [1], [31], [35], [36] and [39]. The authors of [1] propose a state bin transition model to predict the behavior of an aggregate TCL population,

and come up with a linear time-invariant representation of a TCL population that is heterogeneous in *thermal capacitance*. In [31], we use a Markov decision process (MDP) based framework to model the system dynamics of a centralized TCL control scenario using an aggregate population similar to [1].

As more researchers become interested in using TCLs and direct load control for demand response mechanisms, challenges associated with the performance of these mechanisms with respect to the assumed communication capabilities are garnering interest in the research community.

Recent studies, such as [48] and [3], investigate the effects of TCL metering and communications on the performance of load following using a TCL population heterogeneous in all thermal parameters. In particular, the authors develop a Kalman filtering-based approach using the linear model employed in [1] and generate various scenarios to investigate the effects of changing availability on local information on parameter identification, state estimation and control. However, to the best of our knowledge, all of these studies assume that the noise on the aggregate power measurements is white gaussian, and that no other loads are present in the aggregate power measurements obtained.

In this chapter, we first use the MDP-based framework presented in [31, 33] to develop a moving horizon state estimation technique to infer the state of a TCL population that is heterogeneous in all thermal parameters. We use this estimation technique to retrieve the state information based on observations of aggregate power consumption and compare this methodology to benchmark Kalman filtering techniques in the literature. Then, we shift our focus on contextually supervised source separation techniques for energy disaggregation problems recently developed by [52] and adapt our MHSE strategy to incorporate source separation and infer the

underlying TCL population state. To test this strategy, we leverage the dataset used in Chapter 3 (i.e. The Household Energy Survey of the UK). Specifically, we use electricity use data collected from 251 households over the period from May 2010 to July 2011 [4] to create a virtual neighborhood and test our state estimation strategy to investigate state estimation performance, when the aggregate power consumption of all households is measured from a neighborhood level distribution feeder.

4.3 Estimation of TCL States from Aggregate Power Measurements

4.3.1 Individual and Aggregate TCL Models

Individual TCLs are modeled following the approach described in Section 3.1.1. The model representing the temperature dynamics of each TCL is given in (3.1). Similar to [48] and [3], a noise parameter $\epsilon_{i,t}$ is included in the individual TCL model.

We also model the aggregate behavior of individual TCLs based on state bin transitions, as described in Section 3.1.2. For clarity of the state estimation frameworks proposed in this chapter, we include the system of equations describing the aggregate model (i.e. (3.19)) here as well:

$$\begin{aligned} X_{t+1} &= \mathbf{A}^{\text{LTI}} X_t + \mathbf{B}^{\text{LTI}} X_t \circ D_t \\ Y_t &= \mathbf{C} X_t \end{aligned} \tag{4.1}$$

\mathbf{A}^{LTI} and \mathbf{B}^{LTI} are $\{2N^{\text{int}} \times 2N^{\text{int}}\}$ matrices shown in (3.20) and (3.21), respectively.

4.3.2 The Moving Horizon State Estimator

In this chapter, we use the model given in (4.1) to develop the MHSE via an optimization routine that is based on observations of aggregated power Y_t and the decisions D_t .

The MHSE uses the state transition function given in (4.1), where X_t is the state vector formed by $x_{i,t}^{ON}$ and $x_{i,t}^{OFF}$, $\forall i \in [1, N]$, and D_t is the decision/control action at time t as defined in (3.5) and (3.7), respectively. Assume we have an observation Y_t of the aggregate power consumption of the heterogeneous TCL population and a matrix C that relates \hat{X}_t , an estimate of the current state, to an output \hat{Y}_t :

$$\hat{Y}_t = C\hat{X}_t \quad (4.2)$$

We assume that the rated power of each appliance, P_{rated} , is randomly drawn from a uniform distribution within a range of known values. Therefore, for a certain number of appliances N_{app} , the C matrix is formed as:

$$C = N_{app}\hat{P}_{mean}[\underbrace{0 \dots 0}_{N_{int}} \underbrace{1 \dots 1}_{N_{int}}] \quad (4.3)$$

where:

$$\hat{P}_{mean} = \frac{\sum_{i=1}^{N_{app}} P_{rated,i}}{N_{app}} \quad (4.4)$$

Hence, for a time horizon T , we can use the following optimization routine to obtain an estimate of the current state, X_t , by using the observations Y_t (e.g., the aggregate

power consumption) and D_t , the decision vector sent to the plant:

$$\begin{aligned}
 & \underset{\hat{X}_j}{\text{minimize}} && \sum_{t-T+1}^t (Y_j - \hat{Y}_j)^2 \\
 & \text{subject to} && \hat{X}_j = \mathbf{A}^{\text{LTI}} X_{j-1} + \mathbf{B}^{\text{LTI}} X_{j-1} \circ D_{j-1} \\
 & && \hat{x}_{i,j}^{ON} \geq 0, \\
 & && \hat{x}_{i,j}^{OFF} \geq 0, \\
 & && \hat{X}_j \vec{1} = 1, \\
 & \text{where} && j \in [t - T + 1, t], i \in [1, N].
 \end{aligned}$$

As can be seen in (4.3.2), the proposed MHSE uses a cost function that only penalizes the distance between the expected and measured aggregate power consumption value. A more generic moving horizon estimation formulation can be obtained by introducing the *arrival cost* term and incorporating the system dynamics given in (4.1) in the cost function [53, 54, 55].

4.3.3 Kalman Filtering

In order to implement the Kalman filtering approach described in [48] and [3], a linear model of the system is obtained using the Markov chain model introduced in [1]. Using the state-bin transition model depicted in Figure 3.1, the natural dynamics of the system are described via a Markov transition matrix, whose elements denote the transition probability in between status and interval pairs. The linear system can be expressed as initially introduced in (3.9):

$$\begin{aligned}
 X_{t+1} &= \mathbf{A}^{\text{bLTI}} X_t + \mathbf{B}^{\text{bLTI}} X_t \circ D_t \\
 Y_t &= \mathbf{C} X_t
 \end{aligned} \tag{4.5}$$

where $X_t \circ D_t$ can be referred to as u_t . The input vector u is used to control the amount of probability mass transferred within each temperature interval between different appliance statuses. The \mathbf{A}^{bLTI} matrix can be obtained by calculating the maximum likelihood estimate of transition probabilities between different temperature interval and status pairs (similar to [1]) using a simulated population of TCLs with predetermined characteristics. The $\{2N_{int} \times 2N_{int}\}$ matrix \mathbf{B}^{bLTI} is defined using identity matrices of size N_{int} , $\mathbb{1}_{N_{int}}$, as follows, to keep the next state vector a valid probability mass distribution:

$$\mathbf{B}^{\text{bLTI}} = \begin{bmatrix} -\mathbb{1}_{N_{int}} & \mathbb{1}_{N_{int}} \\ \mathbb{1}_{N_{int}} & -\mathbb{1}_{N_{int}} \end{bmatrix} \quad (4.6)$$

In [48] and [3] the authors ensure that each temperature interval receives a single switching probability since bi-directional actions do not provide any benefits in terms of realizability of the desired change in the state and they increase the overall number of switchings. The C matrix can be defined in the same way as in (4.3).

To implement the Kalman filter, we assume perfect measurements and the following linear model for system dynamics:

$$\begin{aligned} X_{t+1} &= \mathbf{A}^{\text{bLTI}} X_t + \mathbf{B}^{\text{bLTI}} u_t + \mathbf{B}_\omega \omega_t \\ Y_t &= \mathbf{C} X_t + v_t \end{aligned} \quad (4.7)$$

where ω_t is the process noise vector. The process noise sequence is assumed to be white, independent random processes with normal probability distribution [56]:

$$p(\omega_t) \sim N(0, Q)$$

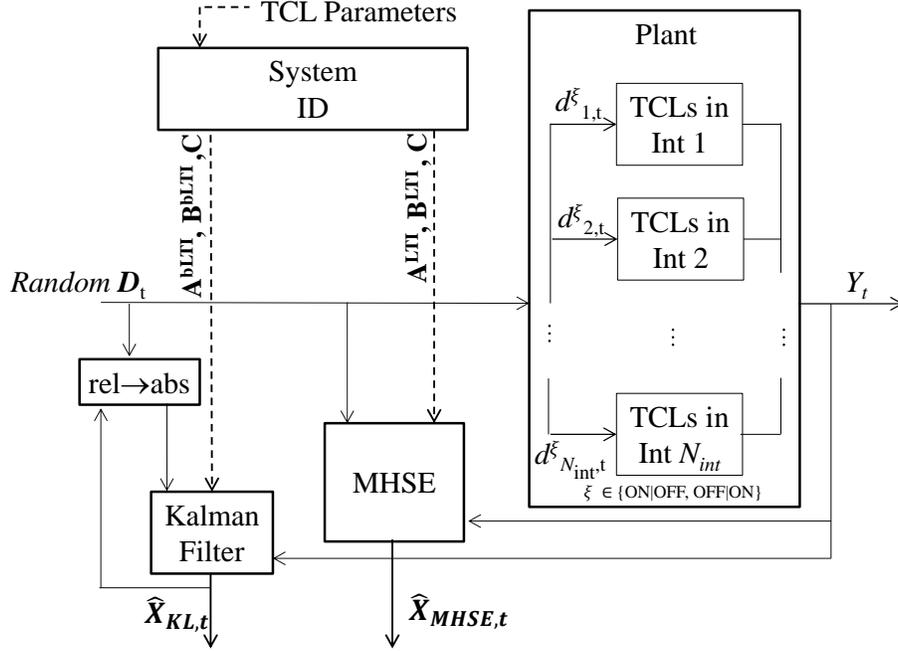


Figure 4.1: Estimation Model

The process covariance matrix Q is computed with the residuals between the state values predicted by the model and the real state values. Since a perfect Q matrix is hard to obtain, we do not update the Q matrix in each estimation time step. Instead, we compute Q for the given system parameterization in the case study and assume it is constant throughout the estimation process.

4.4 Case Studies

To test the performance of the proposed state estimator, we create 2 different case studies. In the first case study, we test the performance of the MHSE for varying window size, T . In the second case study, we test the performance of the Kalman filter and the MHSE estimator under different switching conditions.

Both estimation routines are depicted in Figure 4.1, where $\hat{X}_{KF,t}$ and $\hat{X}_{MHSE,t}$ are used to differentiate between the estimates of X_t obtained from the Kalman

filter and the MHSE, respectively. As it can be seen in Figure 4.1, both estimators are able to receive the aggregate power consumption of the TCL population and random control actions, D_t taken at each time step. For the Kalman filtering technique, we are using the estimated states, $\hat{X}_{KF,t}$ to transform between the switching probabilities that are randomly generated and the control signal u_t described in Section 4.3.3. In addition, the estimates of the Kalman filter are extrapolated at each time step to represent a valid probability mass distribution.

To quantify the information lost when the real state X_t is approximated by the estimated state \hat{X}_t , the Kullback-Leibler (KL) divergence $D_{KL}(X_t||\hat{X}_t)$ [57] is used. KL divergence of X_t from \hat{X}_t is defined as:

$$D_{KL}(X_t||\hat{X}_t) = \sum_i^{2N} \ln \left(\frac{x_{i,t}}{\hat{x}_{i,t}} \right) x_{i,t} \quad (4.8)$$

For each simulation run, we calculate the average KL divergence of X_t from \hat{X}_t over time as:

$$\hat{D}_{KL}(X_t||\hat{X}_t) = \frac{\sum_{t=0}^{T_{tot}} D_{KL}(X_t||\hat{X}_t)}{T_{tot}} \quad (4.9)$$

4.4.1 Case Study I

In this case study, our goal is to understand the effect of the time horizon, T , on the proposed state estimation technique that is using (4.1). To achieve that, we run 10 simulations for each of the 6 different values of the window size, T : 10, 20, 30, 40, 50 and 60 minutes. We identify the system parameters of the MHSE using the simulation parameters given in Table 4.1. During these simulations, there are no limitations on the switching probabilities $d_{k,t}^\xi$ that are sent to the loads. The 95% confidence interval for the mean KL divergence $\hat{D}_{KL}(X_t||\hat{X}_t)$, calculated using the results obtained from simulations with different time horizon T values, is depicted

in Figure 4.2.

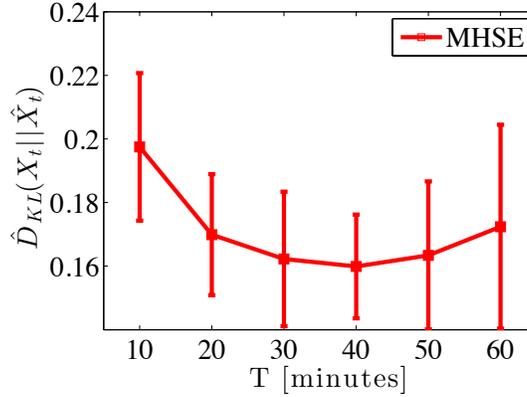


Figure 4.2: Mean KL Divergence, $\hat{D}_{KL}(X_t||\hat{X}_t)$, for Different Time Horizons, T

As seen in Figure 4.2, the optimal estimation results are obtained when T is 40 minutes. When a shorter time horizon T is used, the mean KL divergence $\hat{D}_{KL}(X_t||\hat{X}_t)$ increases, which can be explained by the lack of information available to the estimator regarding past decisions and aggregated power consumption. Although there is more information available to the estimator, the optimization performance decreases for time horizons longer than 40 minutes as well. This can be explained by two factors: (i) the increasing variability in the aggregate power consumption values due to rated power value selection from a uniform distribution, and (ii) the noise process included in the individual temperature dynamics of loads.

4.4.2 Case Study II

In this case study, our goal is to test the performance of the Kalman filter and the MHSE under different switching conditions. We identify the system parameters for the Kalman filtering and the MHSE using the parameters given in Table 4.1. We set the R^{mes} value to zero, and to obtain Q we use the methodology described in Section 4.3.3.

To observe the effect of different values of switching probabilities sent to each temperature interval, we create different scenarios where these switching probabilities, $d_{i,t}^{\xi}$, are bound by the forcing parameter, f , at any given time t such that:

$$\forall d_{i,t}^{\xi} \in [1, N], d_{i,t}^{\xi} \leq f \quad (4.10)$$

This means that the value f is a limit to the values that can be randomly assigned to the elements of D_t . To visualize the effect of the forcing parameter to the system, we use scaled images of the state vectors at each time step.

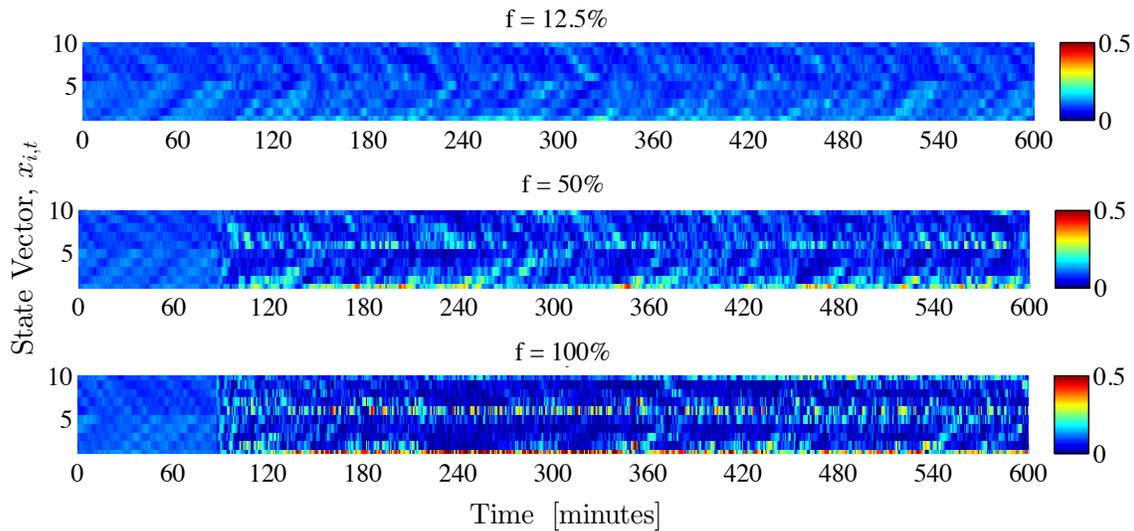


Figure 4.3: State Vector Elements, $x_{i,t}$ vs. Simulation Duration T_{tot} , for f : 12.5%, 50% and 100%

Figure 4.3 shows these images for 3 different forcing parameters f : 12.5%, 50% and 100%. Each sub-image constitutes $\{2N \times T_{tot}\}$ pixels and each pixel is color-coded to show red when the probability mass corresponding to that temperature interval and status pair is above 0.5 and blue when it is 0.

As seen in Figure 4.3, increasing the forcing parameter yields a deviation from the steady state distribution of the probability mass represented by the state vector

<i>Parameter</i>	<i>Definition</i>	<i>Value Case Study I</i>	<i>Value Case Study II</i>
θ_{set}	temperature setpoint	20°C	
δ	temperature deadband width	0.5°C	
θ_a	ambient temperature	32°C	
C	capacitance of appliances	8 – 12 kWh/°C	
R	resistance of all appliances	1.5 – 2.5°C/kW	
h	time step	1 minute	
P_{rated}	rated power of appliances	10 – 18 kW	
$\epsilon_{i,t}$	noise process for individual TCLs	$N(0, 0.01)$	
N	number of temperature intervals	5	
N^{load}	number of appliances	500	
T_{tot}	simulation duration where random actions are received	600 minutes	
T	time horizon	10, 20, 30, 40, 50, 60 minutes	40 minutes
f	forcing parameter	N/A	12.5, 25, 50, 75, 100%

Table 4.1: Simulation Parameters for State Estimation

X_t . In particular, increasing the probability to toggle appliances ON/OFF up to 100% yields accumulations of appliances around the marginal temperature intervals.

To evaluate the performance of both estimators under different forcing conditions, f : 12.5%, 25%, 50%, 75% and 100%, 10 simulations are run with different sets of parameters given in Table 4.1. Each simulation is run for a 720-minute period and random actions are sent to the plant during the last 600 minutes (i.e. $T_{tot} = 600$ minutes).

Figure 4.4 depicts the 95% confidence interval for mean KL divergence, $\hat{D}_{KL}(X_t||\hat{X}_t)$ taken over the estimation period against different forcing values. As seen in Fig-

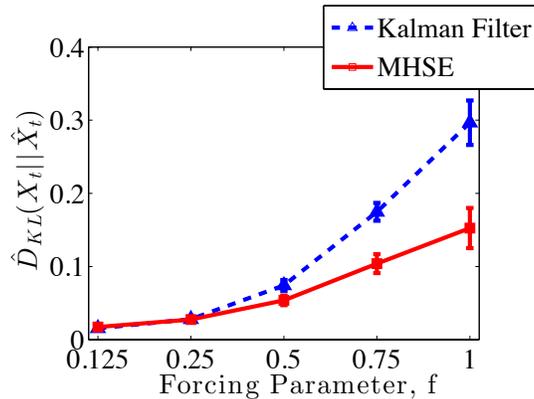


Figure 4.4: Mean KL Divergence, $\hat{D}_{KL}(X_t || \hat{X}_t)$, for Each Forcing Value, f

Figure 4.4, in both scenarios, the KL divergence values obtained from both techniques are similar for low forcing values. Simply put, when the expected number of switchings is low, both estimators perform similarly. However, as forcing increases, the MHSE provides better estimates in comparison to Kalman filtering.

4.5 Using Contextually Supervised Source Separation for State Estimation

In this section, we relax the assumptions made about the availability of aggregate power measurements of the TCL population, and assume that only neighborhood-level aggregate power measurements are available to the main controller and a portion of the neighborhood-level aggregate power measurements consist of loads other than TCLs. We propose a state estimation framework that estimates the states of a TCL population as it separates the aggregate load consumption of the TCL population from neighborhood level aggregate measurements using contextually supervised source separation (CSSS) techniques [52]. We refer to the proposed state estimation framework as the *contextually supervised state estimator* (CSSE).

4.6 Methodology

Assuming the main controller observes the aggregate load signal Y^{agg} , we are interested in the source separation problem where the aggregate power consumption Y^{agg} is composed of L number of loads $i = \{1, \dots, L\}$ Y^i such that for any time t :

$$Y_t^{agg} = \sum_{i=1}^L Y_t^i \quad (4.11)$$

For the most recent historical time window of size T at time t , we form the matrix $\mathbf{Y}_{[t-T+1,t]}$ whose columns are the unobserved power consumptions of each load i :

$$\mathbf{Y}_{[t-T+1,t]} \in R^{T \times L} = \begin{bmatrix} Y_{t-T+1}^1 & Y_{t-T+1}^2 & \dots & Y_{t-T+1}^L \\ \vdots & \vdots & & \vdots \\ Y_t^1 & Y_t^2 & \dots & Y_t^L \end{bmatrix} \quad (4.12)$$

In [52], the authors show that, if each load Y^i can be represented by $\Phi^i \theta^i$, where Φ^i terms are component-specific bases (features) and θ^i are signal coefficients, then obtaining the individual loads in the time period $[t - T + 1, t]$, $Y_{[t-T+1,t]}^i$ can be cast as the following optimization problem:

$$\begin{aligned} & \underset{Y_{[t-T+1,t]}^i, \theta^i}{\text{minimize}} && \sum_{i=1}^L \{ \alpha^i l^i(Y_{[t-T+1,t]}^i, \Phi^i \theta^i) + \beta^i g^i(Y_{[t-T+1,t]}^i) + \gamma^i h^i(\theta^i) \} \\ & \text{subject to} && Y_{[t-T+1,t]}^{agg} = \sum_{i=1}^L Y_{[t-T+1,t]}^i \end{aligned} \quad (4.13)$$

where l^i , g^i and h^i are convex load dependent functions through which an overall convex cost function is constructed. l^i is a loss function which penalizes deviation

between the reconstructed load and its linear representation. g^i is a direct function of the reconstructed load and can be used to reflect additional contextual knowledge regarding the load behavior, and h^i can be used as a regularization term. The specific bases Φ^i are selected based on contextual knowledge about the characteristics of the load i .

To incorporate the model presented in (3.19) to the source separation strategy presented in (4.13), we propose the following optimization strategy in which the aggregate power consumption of TCLs are considered as an unobserved single load type i . We refer to the aggregate power consumption of TCLs as Y_t^{TCL} . Similar to the MHSE, we incorporate the system dynamics as a constraint, and make use of the linear relationship of Y_t^{TCL} and X_t to ensure that the linear representation requirement of the optimization framework given in 4.13 is fulfilled. Then, the proposed CSSE strategy is given as follows:

$$\begin{aligned}
& \underset{Y_{[t-T+1,t]}^i, \theta, X}{\text{minimize}} && \sum_{i=1}^L \{ \alpha^i l^i(Y_{[t-T+1,t]}^i, \Phi^i \theta^i) + \beta^i g^i(Y_{[t-T+1,t]}^i) + \gamma^i h^i(\theta^i) \} \\
& \text{subject to} && Y_j^{agg} = Y_j^{TCL} + \sum_{i=1}^{L-1} Y_j^i \\
& && X_{j+1} = \mathbf{A}^{\text{LTI}} X_j + \mathbf{B}^{\text{LTI}} X_j \circ D_j \\
& && Y_j^{TCL} = \mathbf{C} X_j \\
& && \hat{x}_{k,j}^{ON} \geq 0, \\
& && \hat{x}_{k,j}^{OFF} \geq 0, \\
& && \hat{X}_j \vec{\mathbf{1}} = 1, \\
& \text{where} && j \in [t - T + 1, t], k \in [1, N].
\end{aligned} \tag{4.14}$$

Choosing the functions l^i , g^i , h^i and the number of loads L is challenging without using prior knowledge of individual loads that could be in the aggregate signal. As suggested in [52], we think that a natural choice for l^i is the l_1 norm between the reconstructed load and its linear model. However, depending on the time horizon T , l_2 norm can be used with an h^i function that regularizes the corresponding signal coefficients Θ^i . Based on our results of the MHSE, we expect that for $T \gg h$, where T is the time horizon and h is the simulation time step, the estimation performance will decay. However, as demonstrated in [52], the proposed CSSS strategy is expected to perform better for $T \gg h$. Hence, we expect that the choice of l^i , g^i and h^i will play a significant role in our state estimation strategy. We also expect that other single appliances present in the aggregate signal will rapidly alternate between ON/OFF states. Hence, the l_1 norm of $ZY_{[t-T+1,t]}^i$, where Z is the linear difference operator that subtracts $Y_k^i - Y_{k-1}^i \forall k \in [t - T + 1, t]$, could be useful when representing them in the optimization cost function.

In addition, the parameters α^i capture the *relative* importance of deviations of the individual loads from their linear representations (i.e. their values are inversely proportional to the tolerance of the algorithm to deviate from the proposed linear model). Similarly, the parameters β^i and γ^i serve as weighting factors while forming the aggregate cost function minimized in the strategy described in (4.14). Hence, we expect that the proposed strategy requires tuning for different time window T , control period K and N_{int} parameters, as well as different functions l^i , g^i and h^i .

In the remainder of this chapter, we aim to show that our proposed strategy has the potential to be used as a state estimation strategy, and we investigate the impacts on CSSE performance of different assumptions on the availability of information on the TCLs and other loads present in the aggregate measurement.

4.7 Dataset and Daily Load Simulation

In order to create a test bed to evaluate the proposed CSSE strategy, we use the UK household energy survey described in Section 3.2.1. The Household Energy Survey of the UK includes data on electrical power demand of 251 households over the period from May 2010 to July 2011. It includes multiple datasets with different numbers of appliances, households, collection durations, and sampling rates [4].

Similar to Chapter 3, we use the energy demand measurements collected every 2 minutes from all of the appliances over approximately one month within the May 2010 to July 2011 period. We separate the data into individual days. Even though most of the data from individual households was collected on separate days, we combine daily household data to create a neighborhood-level virtual aggregation point. This means that the household level data is not collected on the same day for all of the households. Hence, additional daily variations are present in the individual household loads used for the same day (e.g., different daily external temperature profiles, day of week etc.). The energy measurements are converted into power measurements assuming constant power draw within each sampling period.

Following daily data extraction, we identify the household refrigeration units (HRUs) in each household and remove them from the dataset. To create a TCL population that can be controlled by the main controller, we use the individual load modeling strategy described in Section 3.1.1 to replace the HRUs removed from the dataset. In particular, we simulate 373 HRUs in a time synchronized manner using the same sampling period as the non-TCL loads ($h=2$ minutes). During simulation, these HRUs are controlled by the main controller via random actions with forcing parameter f .

The parameters used to simulate TCLs and their availability to the estimator

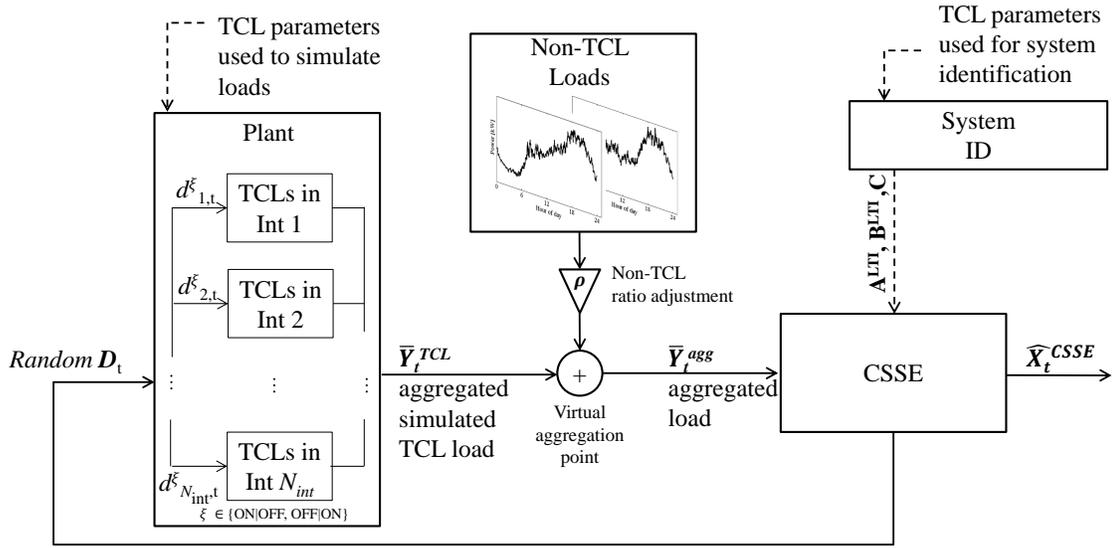


Figure 4.5: CSSE Model

play an important role in the overall performance of the CSSE. For example, if the parameters used for system identification for the estimator are the same as those used to simulate individual loads, the estimator will be biased towards the simulated population. In the next section, we introduce case studies to understand the impact of the availability of information to the estimator and the sensitivity of CSSE's performance to various input parameters.

4.8 Case Studies

In this section, we first introduce the common assumptions made in both of the case studies and describe the performance metrics used to evaluate the proposed CSSE. Then, we focus on individual case studies and elaborate on case-specific assumptions.

In this study, similar to [52], we assume that the aggregate load $Y_{[t-T+1,t]}^{agg}$ is composed of 3 different load types: the population of HRUs, a base load that is time-of-day dependent, and a group of appliances that rapidly switch between ON and

Individual Load, i	Features	l_i	g_i	h_i
$Y_{[t-T+1,t]}^{HRU}$	$ALTI$, $BLTI$ and C	$\ Y_{[t-T+1,t]}^{HRU} - CX_{[t-T+1,t]}\ _1^1$	None	None
$Y_{[t-T+1,t]}^{base}$	2 minute interval of the day	$\ Y_{[t-T+1,t]}^{base} - \Phi^{base}\theta^{base}\ _2^2$	$\ ZY_t^{base}\ _2^2$	$\ \theta^{base}\ _2^2$
$Y_{[t-T+1,t]}^{app}$	None	$\ Y_{[t-T+1,t]}^{app}\ _1^1$	$\ ZY_{[t-T+1,t]}^{app}\ _1^1$	None

Table 4.2: Function selection for case studies

OFF states. To capture the time-of-day dependency of $Y_{[t-T+1,t]}^{base}$, we divide the day into 720 time periods which translates into 2 minute intervals. We propose the l_i , g_i and h_i functions given in Table 4.2 for each load $Y_{[t-T+1,t]}^{HRU}$, $Y_{[t-T+1,t]}^{base}$ and $Y_{[t-T+1,t]}^{app}$. In this study, we focus on identifying the underlying TCL states; hence, unlike [52], we focus solely on the state estimation performance of the CSSE. We acknowledge that the state estimation performance affects the overall source separation performance; however, for the purposes of this chapter, we include $Y_{[t-T+1,t]}^{base}$ to represent a time-of-day dependent base load, and $Y_{[t-T+1,t]}^{app}$ to absorb the variation on the base load. We neither divide the non-TCL loads into validation sets for Y^{base} or Y^{app} , nor report performance results for these loads, since additional assumptions are required to cluster the non-TCL loads into these two categories to form such a validation set.

To understand the effect of g^i and h^i functions on the state estimation performance, we look at 2 separate simulations with the same input parameters. In the first scenario we only include the l^i functions given in Table 4.2 in the cost function. In the second scenario, we also use the additional g^i and h^i functions, as described in Table 4.2.

To test if the aggregate power measurements obtained at the virtual aggregation point provide any benefits on the state estimation over a simple prediction scenario, we randomly sample 100,000 different initial states X_0 as valid probability vectors

from the sample space \mathbf{S} at the beginning of the estimation time period K . We then predict the states using the trained system model. The probability vector sampling problem is identical to sampling from a standard simplex $\Delta^{2N_{int}-1}$ which is given as [58]:

$$\Delta^{2N_{int}-1} = \left\{ (x_0, \dots, x_{2N_{int}-1}) \in R^{2N_{int}} \mid \sum_{i=0}^{2N_{int}-1} x_i = 1 \text{ and } x_i \geq 0 \text{ for all } i \right\} \quad (4.15)$$

An identical sampling problem can be cast using the symmetrical Dirichlet distribution, where $\Gamma(s)$ denotes the gamma function and the parameter vector α is transformed into a single variable $\alpha = 1$ [59]. The probability density function then becomes:

$$f(x_1, \dots, x_{K-1}; \alpha) = \frac{\Gamma(\alpha K)}{\Gamma(\alpha)^K} \prod_{i=1}^K x_i^{\alpha-1} \quad (4.16)$$

To quantify and compare the performance of the CSSE in both of the case studies and the simple prediction scenario, we use two separate metrics. We calculate the root mean squared error (*RMSE*) and the R^2 metric between the elements of the predicted and the real state vectors over the control period K . For all elements i of the state vector X_t , we calculate the R^2 value using the real state values $x_{i,k}$, the average real values $\bar{x}_{i,k}$, and the estimated values $\hat{x}_{i,k}$ over the estimation period K using the following equation:

$$R_i^2 = 1 - \frac{\sum_{k=1}^K (x_{i,k} - \hat{x}_{i,k})^2}{\sum_{k=1}^K (x_{i,k} - \bar{x}_i)^2} \quad (4.17)$$

Table 4.3: Individual TCL parameters obtained from [5] for case study I

	R_i ($^{\circ}\text{C}$)	C_i ($\text{kWh}/^{\circ}\text{C}$)	P_i^r (kW)	θ_i^a ($^{\circ}\text{C}$)	δ_i ($^{\circ}\text{C}$)	$\theta_{i,set}$ ($^{\circ}\text{C}$)
Distribution	Uniform	Uniform	Uniform	Constant	Uniform	Uniform
Range	80-100	0.4-0.8	0.2-1.0	20	1-2	1.7-3.3

similarly, for all i we calculate the $RMSE$ value using the following equation:

$$RMSE_i = \sqrt{\frac{\sum_{k=1}^K (x_{i,k} - \hat{x}_{i,k})^2}{K}} \quad (4.18)$$

Finally, to obtain an overall state estimation performance score, we calculate the average R_i^2 and $RMSE_i$ among all bins i . We refer to these values as \overline{R}^2 and \overline{RMSE} throughout this section. We use two metrics to understand the improvement over the simple prediction scenario both in absolute and relative terms. The $RMSE$ is an absolute deviation measure between the estimated and real states. Simply put, it is an indicator of how close the estimated states are to the real states. However, the R^2 metric is a relative measure of *fit*. It can be interpreted as the proportional improvement in estimation, compared to the deviation of real states from their mean values.

In the next section, we describe the individual case studies and introduce case-specific assumptions when necessary.

4.8.1 Case Study I

In this case study, we assume that the parameters used to simulate the TCLs are available to the estimator and that the system identification is done on a separate population simulated using identical parameters. We obtain these parameters from [5], as included previously in Table 3.3. For clarity, we include the same table for the reader in Table 4.3.

Table 4.4: Input parameters for case studies I and II.

<i>Parameter</i>	<i>Description</i>	<i>Case Study I</i>		<i>Case Study II</i>
		<i>Scenario a</i>	<i>Scenario b</i>	
N^{int}	number of temperature intervals	6		
K	estimation duration	60 minutes		
T	moving-window size	10,60,120 minutes		
f	forcing parameter	10%, 50%, 100%	100%	
μ	non-TCL to TCL ratio	1	1, 3, 5	

After obtaining the parameters, we run the CSSE with different input parameters under two different scenarios. For the first scenario, we run the estimator under varying forcing parameters with all other parameters kept constant to understand the effects of forcing in the estimation performance. In the second scenario, we keep the forcing value constant (i.e. $f = 100\%$) and analyze the effect of changing the μ parameter on the performance of the estimator. In the dataset, we observe that the maximum aggregate TCL load is approximately 20% of the non-TCL load. In our analysis, we investigate cases with $\mu = 1$, where the maximum TCL population load is equal to the maximum non-TCL load, $\mu = 3$, where the maximum TCL population load is approximately 33% of the maximum non-TCL load, and finally $\mu = 5$, which represents a realistic case where the maximum TCL population load is 20% of the maximum non-TCL load. The input parameters used in both scenarios under this case study are given in Table 4.4.

Figure 4.6 shows the results of scenario *a*. The first row of graphs shows the performance of the CSSE in terms of \overline{R}^2 , and the second row of graphs shows the same

results in terms of \overline{RMSE} . The first three pairs of columns include the results for the varying forcing parameters. For each forcing parameter, the CSSE's performance is presented with respect to different time horizons T on the left graphs. As described in Section 4.8, two simulation runs with identical input parameters but different cost functions are made. The results of the first simulation run with a cost function composed of only the l^i terms are labeled as $\overline{R}_{CSSE,residuals}^2$ and $\overline{RMSE}_{CSSE,residuals}$, respectively. Similarly, the results of the simulation run using all of the l^i , g^i and h^i functions given in Table 4.2 are labeled as \overline{R}_{CSSE}^2 and \overline{RMSE}_{CSSE} . The whiskers of the CSSE performance plots denote the $\pm\sigma$ range of the respective performance metrics. On the right graphs, a histogram of the \overline{R}^2 and the \overline{RMSE} values obtained from the simple prediction scenario with 100,000 random initial points are given. The red lines denote the expected value of each distribution.

As described earlier, our goal is to provide a proof-of-concept case study that shows the premise of the proposed CSSE. In this regard, we expect to show that the proposed CSSE strategy performs better than the simple prediction case based on random initial conditions. In Figure 4.6, we observe that the simulation runs using only residuals mostly outperform the simulation runs with g^i and h^i terms (i.e. a higher R^2 score and a lower $RMSE$ value). In addition, for all cases where $T \geq 60$ minutes, the CSSE outperforms the simple prediction case, and, for most of the cases we see that the estimation performance increases by increasing time window T , when going from 10 minutes to 60 minutes. However, for $T = 120$ minutes, the performance of the CSSE either remains similar or decreases. We believe that this is because of the strict constraint representing the system dynamics in the optimization framework. Similar to the MHSE, after a certain window size, the performance decays due to increasing model mismatch for longer estimation periods.

For $f = 100\%$, we observe low $\overline{R^2}$ values with low \overline{RMSE} values. We believe that this is due to low forcing in this estimation scenario. The $\overline{R^2}$ distribution obtained for the simple prediction scenario also suggests a similar explanation. We observe that the estimated values do not provide an improvement over the mean estimator as suggested by the $\overline{R^2}$ distribution. As the forcing parameter increases, we see that the \overline{RMSE} value does not change significantly for $T \geq 60$ minutes. However, we observe an improvement in the R^2 value. This indicates that the estimator does a better job in comparison to a mean estimator as perturbations to the system increase.

In Figure 4.7, we depict the results of the simulation runs we did for varying μ values (i.e. scenario b). Our main goal was to inspect the performance of the model as the proportion of TCLs decreased in the aggregate signal observed at the neighborhood-level virtual aggregation point. Similar to Figure 4.6, the first row of graphs depicts the performance of the CSSE in terms of $\overline{R^2}$ and the second row of plots depicts the performance of the CSSE in terms of \overline{RMSE} . The first three columns of graphs show the results for each μ value investigated in this scenario. The last column includes the distribution of $\overline{R^2}$ and \overline{RMSE} for the simple prediction scenario with 100,000 random initial points. It is important to note that since the forcing parameter remains the same, the simple prediction scenario's results are the same for varying μ cases.

In Figure 4.7, we observe that the CSSE outperforms the simple prediction scenario when $T \geq 60$ minutes. We also see that the performance of the CSSE decreases with the increasing proportion of the non-TCL loads in the aggregate signal for all T values. This decrease is apparent when measured with both of the performance metrics used in this section. Furthermore, the variation in the performance values

also increase with the increasing proportion of the non-TCL loads measured at the virtual aggregation point.

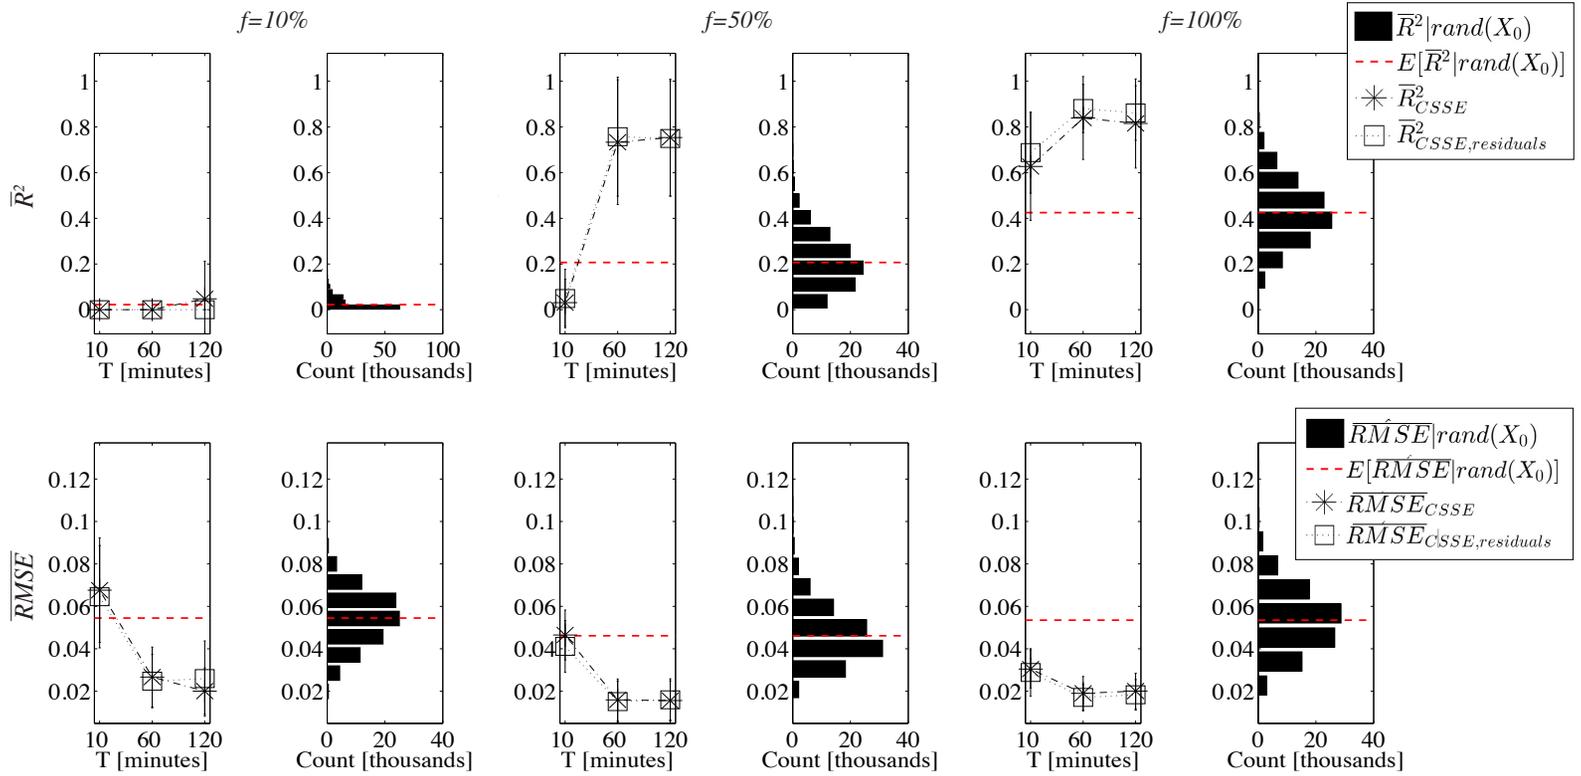


Figure 4.6: Comparison of CSSE performance for varying forcing parameters f (Scenario a)

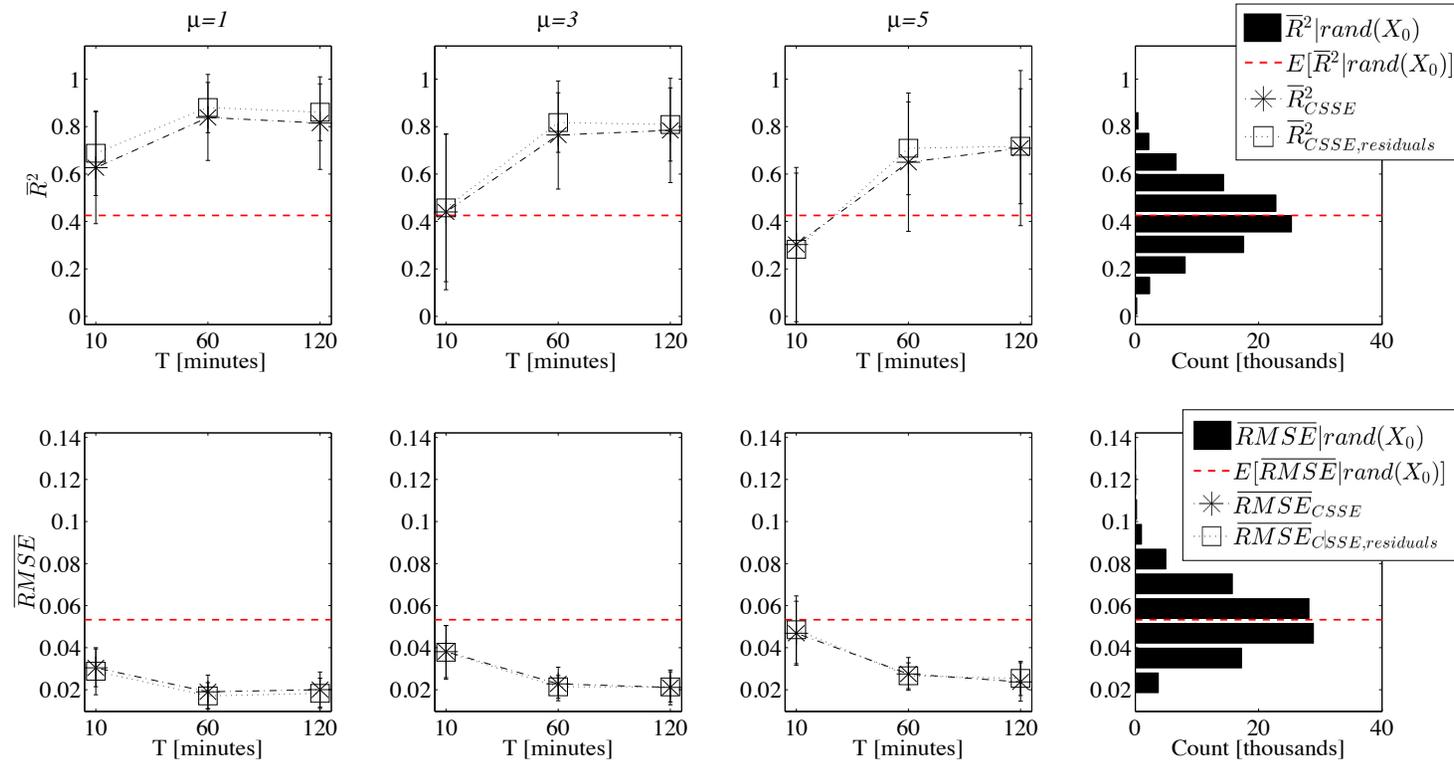


Figure 4.7: Comparison of CSSE performance for varying μ parameters f (Scenario b)

Table 4.5: Individual TCL simulation parameters for case study II

<i>Parameter</i>	<i>Description</i>	<i>Range</i>
θ_i^a	Constant	$20^\circ C$
T_i^{OFF}	Weibull	obtained via Table 3.2
T_i^{ON}	Weibull	obtained via Table 3.2
h	Constant	2 minutes
P_i^r	Uniform	$[0.1 \ 0.2] \text{ kW}$
δ	Constant	$2^\circ C$
θ_i^s	Constant	$2^\circ C$
R	Calculated from T_i^{ON} and T_i^{OFF}	
C	Calculated from T_i^{ON} and T_i^{OFF}	

4.8.2 Case Study II

In case study II, we investigate a scenario where the parameters used to simulate the TCLs are different from the parameters available to the estimator. In particular, we use the thermal parameters proposed in Section 3.2.2 to model the individual TCLs, and we use the thermal properties obtained from [5] as presented in Table 4.3 to perform system identification for the estimator. The motivation behind this is twofold: (i) in a real-world scenario, the information on thermal properties of loads is hard to obtain, and a possible solution might be to leverage existing studies; (ii) the proposed modeling strategy creates a TCL population where the occupant behavior is intrinsically included, as described in Chapter 3. For clarity, we included the parameters initially introduced in Section 3.2.2 in Table 4.5. In this case study, we only focus on varying μ values to investigate the CSSE performance (with respect to the simple prediction scenario with 100,000 initial random conditions) when the model available to the estimator does not fully represent the dynamics of the simulated TCL population. The input parameters are given in Table 4.4. Figure 4.8 shows the results of the performance comparison. The results are depicted following a similar strategy to Figure 4.7 to facilitate comparison.

In Figure 4.8, it is possible to see that for all of the cases, both of the metrics show a poor performance in comparison to the results obtained for scenario b under case study I. When compared to the simple prediction scenario, there is no significant improvement over the expected value of the \overline{RMSE} and \overline{R}^2 distributions.

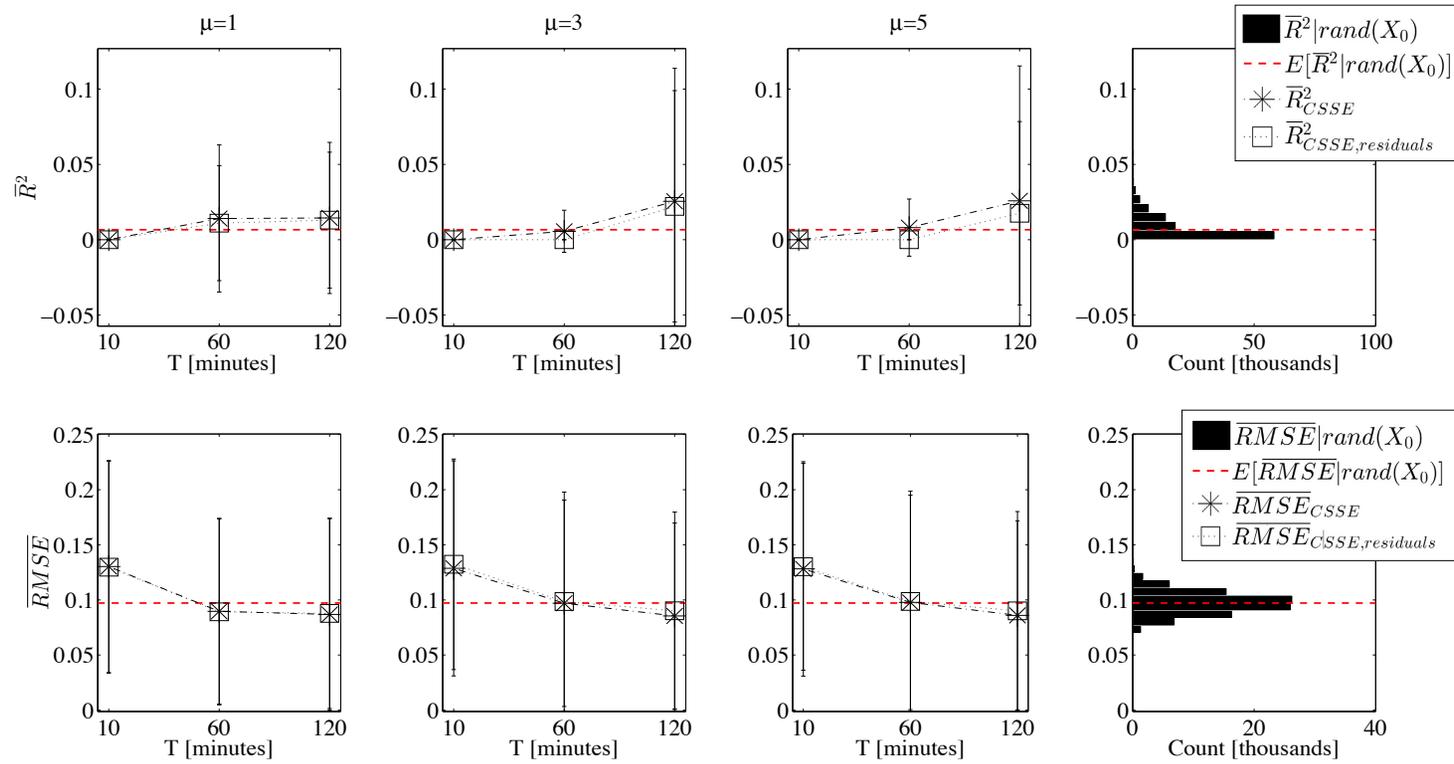


Figure 4.8: Comparison of CSSE performance under the parameters given in case study II.

4.9 Chapter Conclusions

In this chapter, we first developed a moving horizon mean squared error state estimator with constraints to estimate the states of a heterogeneous TCL population from aggregate power measurements of a TCL population. We analyzed the performance of the proposed MHSE under different time horizon T values, and compared the performance of this estimator with a generic Kalman filtering approach based on [3] under different switching conditions imposed by the forcing parameter f . The results show that some improvement is possible in scenarios where the expected switching frequency is higher than 25%. We believe that this increase is due to the nature of the model used for the MHSE; specifically, due to the additional flexibility introduced in the model through changing the way control actions are included in the system dynamics. An increase in the switching frequency is often expected when the controller is trying to keep a constant aggregated power consumption over certain period of time, and when there are abrupt changes in the reference power followed by the TCL population (e.g., load following).

We then relaxed the assumption on the availability of the aggregate power measurements of the TCL population. Specifically, we have introduced the *contextually supervised state estimator* (CSSE) to infer the states of a heterogeneous TCL population from power measurements obtained at a neighborhood-level distribution feeder. As a proof of concept, we showed that the CSSE outperforms a simple prediction strategy in cases where the CSSE has information on the thermal properties of the individual TCLs in the load population. Our results also suggest that, in such cases, the increasing proportion of the non-TCL loads present in the virtual aggregation point decreases the performance of the CSSE; however, it is still possible to see significant improvements over the simple prediction scenario for cases where

the TCL load is approximately 20% of the non-TCL load in the aggregate power signal. Finally, we investigated a case where the CSSE has no information on the thermal properties of the individual TCLs in the load population and it uses the parameters available in the literature to perform system identification. The results of this investigation suggest that the performance of the CSSE is highly dependent on the availability of a reliable model which requires system identification using information on the thermal properties of individual TCLs in the population.

Chapter 5

Estimating the Benefits of Electric Vehicle Smart Charging at Non-Residential Locations

A recent analysis identifying the infrastructure and technology needs to meet California's greenhouse gas (GHG) reduction goals for 2050 shows that the electrification of the transportation system plays a significant role in reaching these goals. In order to achieve the 80% reduction target in electrification, most of the direct fuel uses in buildings, transportation and industrial processes must be electrified. Among these, electrification of transportation yields the largest share of GHG reduction, where 70% of the vehicle miles traveled should be by electrically powered vehicles [27]. Other studies also include forecasts on expected number of EVs in the US by 2020 [60].

The need for rapid growth in the electrification of transportation presents significant challenges as well as opportunities to the operation of today's power system.

This chapter is based on [34].

When considered as inflexible loads, EVs may increase the current electricity demand significantly, intensifying the stress on the electric power system and pushing it closer to its limits. Research suggests that this is the case for low penetrations of EVs [15, 16, 17]. However, when considered as flexible resources, where EV charging is controlled by direct or indirect strategies, EVs promote the reliable operation of the power grid [61, 62, 63], while also providing additional revenue streams that can be used towards the electrification of transportation [15, 62, 22]. This is particularly important considering the expected increase of renewable generation sources in the generation portfolio of many states in the U.S., as smart EV charging may provide the means to balance the intermittency of these resources.

A number of recent studies aim to understand the adaptation needs of the existing operational control mechanisms to realize smart charging, and often propose novel planning and control approaches. These approaches can be grouped into *direct* and *indirect* control approaches [62]. Direct control approaches leverage control strategies in which the control actions are realized without the vehicle owner in the control loop. Often, load aggregations are created to increase the size of the resource so it can offer economic benefits to the aggregator [63, 64]. In [65], for example, the authors propose a direct load control strategy to provide vehicle-to-grid services for 3 different predefined mobility patterns. In [66], the authors conduct a simulation study for 3000 EVs parked at a municipal parking lot and evaluate the real-time performance of a direct control approach, which maximizes the expected state of charge of the EV aggregation in the next time step subject to mobility constraints. In [67], the authors develop an optimal direct control scheme based on global charging costs. The authors compare the proposed direct control scheme to the local scheduler in a simulation environment including 100-400 EVs. The arrival times of

the EVs, the charging periods, and the initial energies of EVs are assumed to have a uniform distribution.

The authors of [64] discuss various services that can be provided by electric vehicles, including peak shaving, regulation, voltage control, and reserves, and many studies have quantified the benefits of smart charging from various stakeholder perspectives [68, 69, 70]. In [64], the authors demonstrate a proof of concept regulation case study. In [68], the authors estimate that smart charging will reduce the daily electricity costs of a plug-in hybrid EV by \$0.23. They also identify daily profits for the individual driver when the charging of the vehicles can be regulated. The economic benefits of fleets that participate in specific markets have also been extensively studied. For example, in [69], 352 vehicles are used to estimate the economic potential of fleets when providing regulation up and down services using historical prices obtained from California Independent System Operator (ISO). In [71], the authors use historical market data and charging data collected from an EV located in a residential household to investigate financial savings and peak demand reduction. The authors conclude that the peak EV demand can be reduced by up to 56%.

In this chapter, we primarily focus on direct control approaches—in particular, centralized smart charging of EV aggregations—and we create a variety of case studies to investigate the potential benefits of smart charging to different stakeholders. To develop these case studies, we use data collected from over 2000 non-residential electric vehicle supply equipments (EVSEs) located throughout 190 zip code regions in Northern California spanning one year. To the best of our knowledge, this is the first study that uses such an extensive dataset on EV charging. First, we analyze over 580,000 charging sessions to investigate the trends in load flexibility and infrastructure use in the dataset. Next, we create virtual aggregation points (VAP)

in which a combination of the EVSEs is assumed to be fed by the same distribution feeder. The VAPs mostly coincide with Pacific Gas and Electric Company's (PG&E) sub-load aggregation points (sub-LAPs). Additional details regarding this relationship is provided in Section 5.1. We introduce a smart charging framework to estimate the benefits of smart EV charging to various stakeholders in each VAP. As an initial case study, we investigate the potential benefits of EV aggregations operated under a single owner, where a time-of-use pricing scheme is used to estimate economic benefits to the owner via shifting load from high cost periods to lower cost periods. Then, we create a case study where EV aggregations are used to decrease their current contribution to the system-level peak load.

The motivation for this study is threefold: (i) Most of the work investigating the potential of smart charging of EVs is based on assumptions made regarding trip and customer characteristics [72, 73, 74, 75]. For example, in [73], the authors use a fleet which includes commuter cars, family cars and taxis with predetermined departure and arrival locations randomly selected from a limited number of alternatives. In [74], the authors use data from driving surveys that reflect the driving behavior of people using internal combustion engine cars [76]. They assume that the driving behavior of an EV owner will be similar to that of an internal combustion engine car owner. The dataset used in this study allows us to extract trip and customer characteristics, hence no such assumptions are needed on these characteristics. (ii) Often, a limited number of vehicles and mobility patterns are used in fleet-based studies to capture the most likely driving scenarios [65, 75, 77]. For example, in [65], the authors develop a proof of concept strategy and show cost benefits for 50 EVs with 3 different pre-defined mobility patterns. Although the exact number of EVs are not available in the dataset used in this study, the number of charging sessions (over

580,000) and the fact that these charging sessions are spread throughout the year ensure that a representative population of non-residential charging is studied. (iii) The individual charging profile of an EV is often represented by a typical constant-voltage, constant-current curve for certain battery chemistries, or more simply by a constant charging power [62, 78, 79, 80, 81]. For example, in [79], the charging power is assumed to be fixed at 4.4kW, whereas in [80], the authors use the charging profile of a typical lithium-ion battery pack obtained from [81]. The dataset used in this study includes time series of power measurements obtained every 15 minutes for each charging session. Hence, no assumptions are made on charging profiles of the vehicles, and individual charging data is available for each charging session.

The remainder of the chapter is organized as follows: Section 5.1 introduces the dataset and discusses the load flexibility and infrastructure use trends obtained from the dataset. Section 5.2 presents the smart charging strategy used in this study. Specifically, it discusses the framework and the underlying assumptions made when estimating the benefits to different stakeholders. Sections 5.3 and 5.4 describe the case studies completed in this research. Finally, Section 5.5 discusses the conclusions and opportunities for future work.

5.1 Dataset

The data used in this study is collected from individual EVSEs located in 16 different sub-LAPs in PG&E's territory for the year of 2013. For each charging session (i.e. from plug-in to departure of an EV), the EVSEs report the start and end period of the charging, the plug-in and departure time stamps, the average power, and the maximum power (measured every 15 minutes), as well as the charging port type, the location (zip-code level), and the non-residential building category. Since

the dataset includes the location information based on zip codes and some zip codes are fed by multiple sub-LAPs, we create virtual aggregation points (VAPs) for the zip codes that are fed by multiple sub-LAPs. This is done by combining the sub-LAPs' identifiers. Table 5.1 presents the final list of VAPs in the dataset and total number of zip code regions forming each of these VAPs, the total number of charging sessions, and the average number of daily charging sessions in each VAP. Figure 5.1 depicts the centroids of the zip code regions forming the considered VAPs.

In this study, we use data from VAPs with an average of 20 or more charging sessions per day. These VAPs are indicated in **bold** in Table 5.1. Figure 5.2 also shows the total number of charging sessions per month for each VAP used in this study. Over the course of 2013, the total number of charging sessions approximately doubles among the VAPs used in this study.

Figure 5.3 shows the combined load profiles of VAPs for the second weeks of January and December. The impact of the growth in charging session is reflected on the daily load profile of the loads. Moreover, the peak non-residential EV load occurs between 9AM and 11AM, and it more than triples from January to December of 2013.

5.1.1 Load Flexibility and Infrastructure Use

To gain further insight into the dataset and to understand the distribution of charging sessions and the use of EVSEs in different regions, we analyze the charging sessions obtained from the VAPs marked in bold Table 5.1. The infrastructure use, I_{use} , in each VAP is represented by the average number of charging sessions $N_{sessions}$ per EVSE and calculated for every business day of 2013. Formally:

$$I_{use} = \frac{N_{sessions}}{N_{EVSE}} \quad (5.1)$$

VAP	Region	# of zip code regions	# of charging sessions	# of charging sessions per day
P2-SB	Peninsula & South Bay	7	207501	568.50
SB	South Bay	21	112250	307.53
SF	San Francisco	30	72996	199.99
P2	Peninsula	17	59252	162.33
EB	East Bay	27	52700	144.38
EB-SB	East Bay & South Bay	6	16902	46.31
NB	North Bay	14	12346	33.82
LP	Los Padres	8	9035	24.75
CC	Central Coast	15	8428	23.09
SA	Sacramento Valley	11	7787	21.33
FG	Geysers	11	7918	21.69
SA-SI	Sacramento V. & Sierra	2	7465	20.45
CC-P2	Central Coast & Peninsula	2	6778	18.57
FG-NB	Geysers & North Bay	4	3845	10.53
F1	Fresno	4	377	1.03
NV	North Valley	1	336	0.92
ST	Stockton	3	244	0.67
FG-NC	Geysers & North Coast	1	246	0.67
SI	Sierra	2	181	0.50
SN	San Joaquin	1	134	0.37
HB	Humboldt	1	101	0.28
P2-SF	Peninsula & San Francisco	1	73	0.20
NC	North Coast	1	15	0.04

Table 5.1: VAPs used in this study

where N_{EVSE} is the number of EVSEs. Figure 5.4 depicts the box plots of the infrastructure use within 2013 for all of the VAPs. For each month of 2013, a box plot is created to represent the distribution of the I_{use} values calculated for every business day of the month. The median value of infrastructure use is marked with a red line in each box plot, and the boundaries of the box depict the 25th and 75th percentiles. The whiskers correspond to the 99th percentiles assuming the distributions per each month are normal. The median infrastructure use increases in all

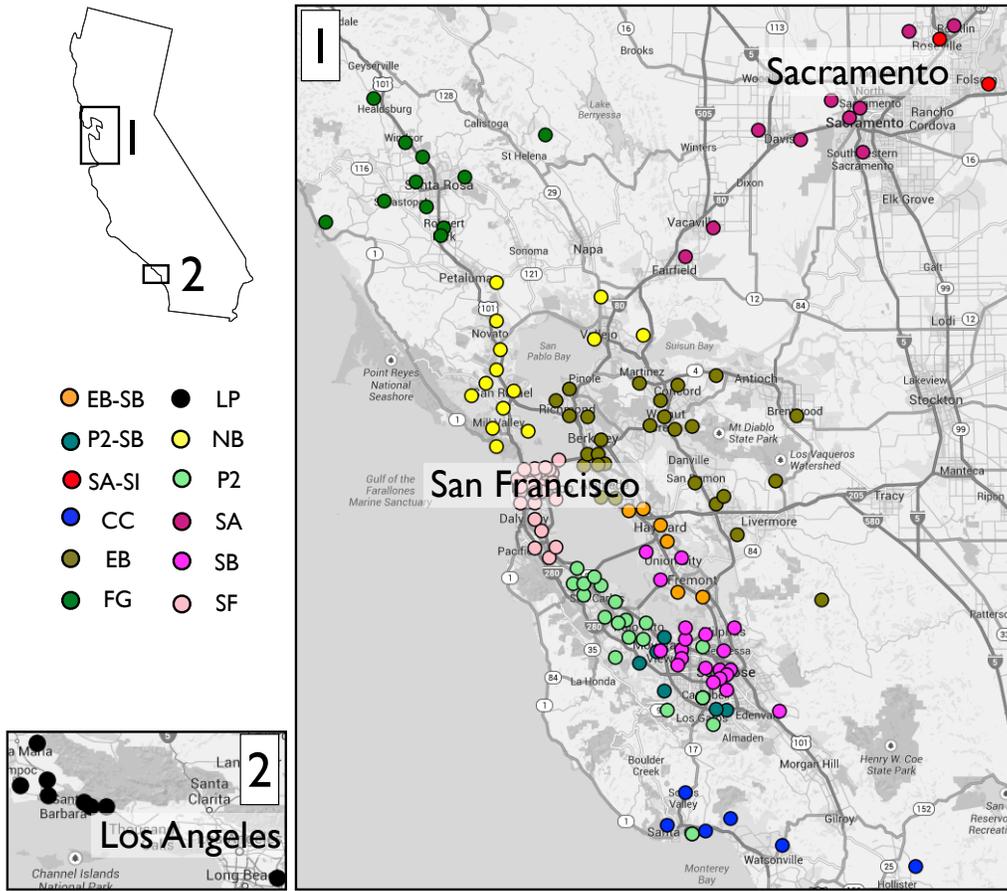


Figure 5.1: Centroids of zip code regions forming the VAPs

VAPs from 1.8 to 2.1 sessions per EVSE from January to December. We believe that this is due to the fact that the demand has increased faster than the number of EVSEs.

In addition to the infrastructure use, we investigate the load flexibility in each VAP. The load flexibility depends on the charging duration d_{charge} and the overall duration of each charging session d_{session} . Formally, we define the load flexibility l_{flex} as the ratio of the duration that a car is plugged but not charging to the overall

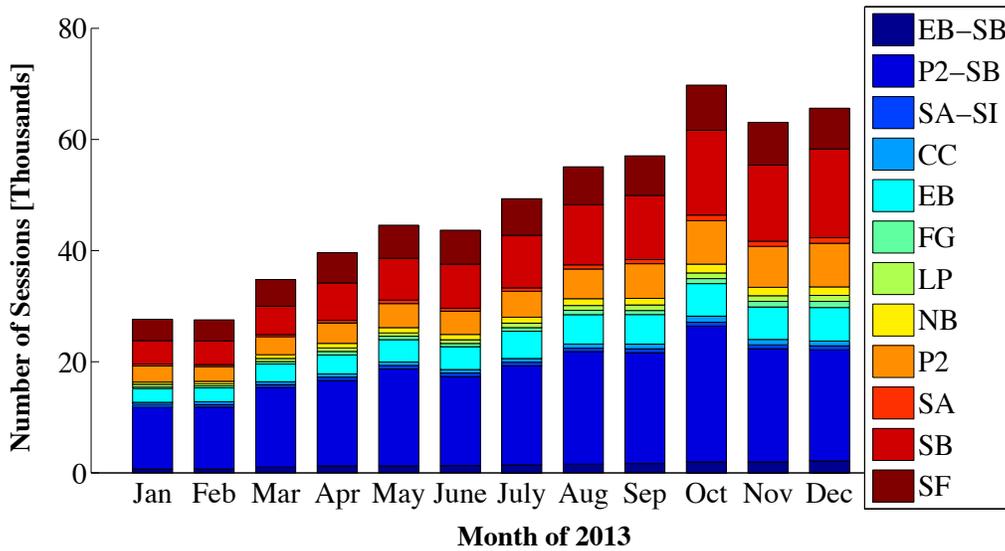


Figure 5.2: Number of sessions per month

session duration:

$$l_{\text{flex}} = \frac{d_{\text{session}} - d_{\text{charge}}}{d_{\text{session}}} \quad (5.2)$$

Figure 5.5 depicts the load flexibility for all VAPs by month. As Figure 5.5 suggests, the load flexibility decreases slowly as the number of charging sessions per EVSE increases. Also, most of the distributions have a slight positive skew. The size of the box representing the 25th and 75th percentiles is also decreasing with time, suggesting an increase in skewness.

The load flexibility metric shows the charging duration relative to the session duration; however, it does not capture when the charging sessions occur. The start and end times of the charging sessions play a key role when estimating the benefits of EV aggregations to the power system. To put these two variables into perspective, we show a histogram of plug-in (i.e. session start) and plug-out (i.e. session end) times in Figures 5.6a and 5.6b, respectively.

As can be seen in Figures 5.6a and 5.6b, most of the charging sessions start

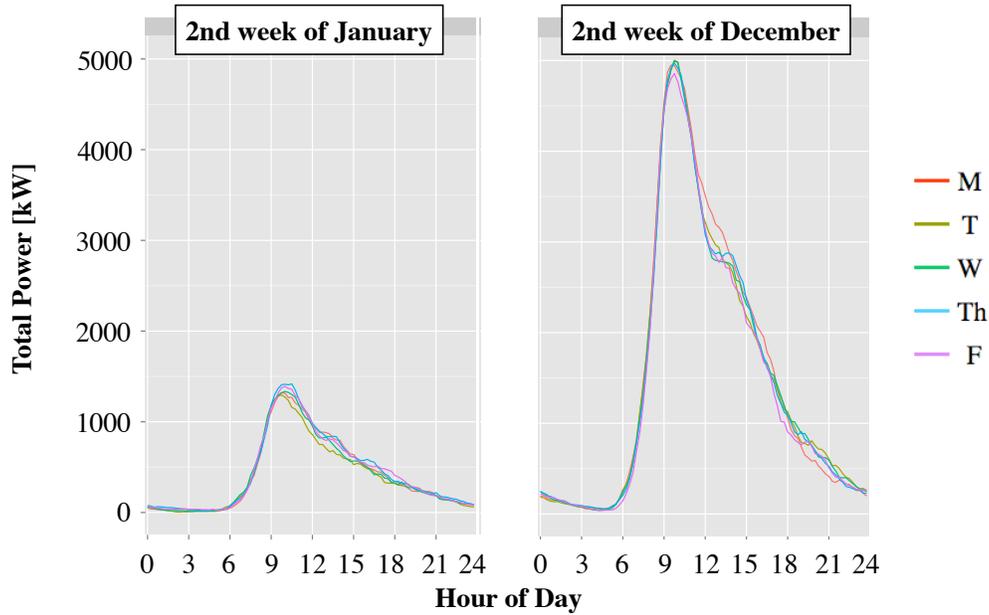


Figure 5.3: Aggregate EV load shapes for January and December for all the VAPs

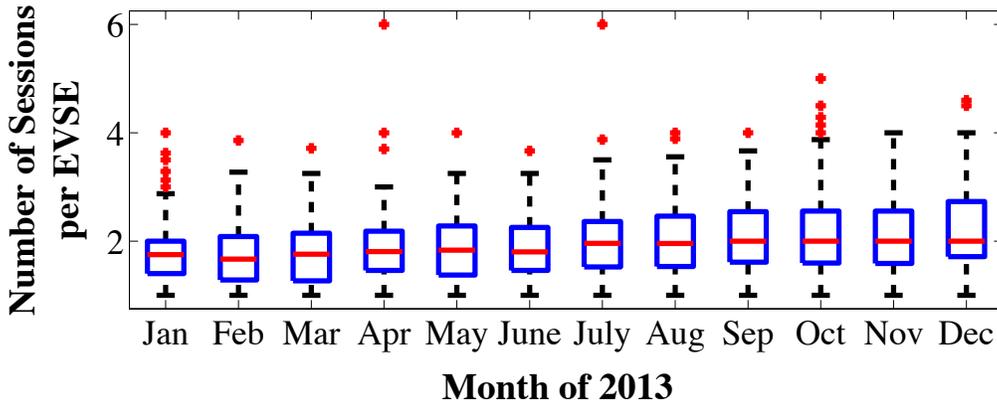


Figure 5.4: Average number of sessions per unique EVSE per day

within the 7AM-10AM period and often end within the 5PM-7PM period. Considering these loads are currently uncontrolled (i.e. they immediately start charging when they are plugged in), they coincide with the typical working hours of a non-residential location. These figures suggest that most of the employees or customers plug in their vehicles to the EVSEs in the morning. Some unplug their vehicles around noon and most unplug their vehicles between 4PM and 7PM.

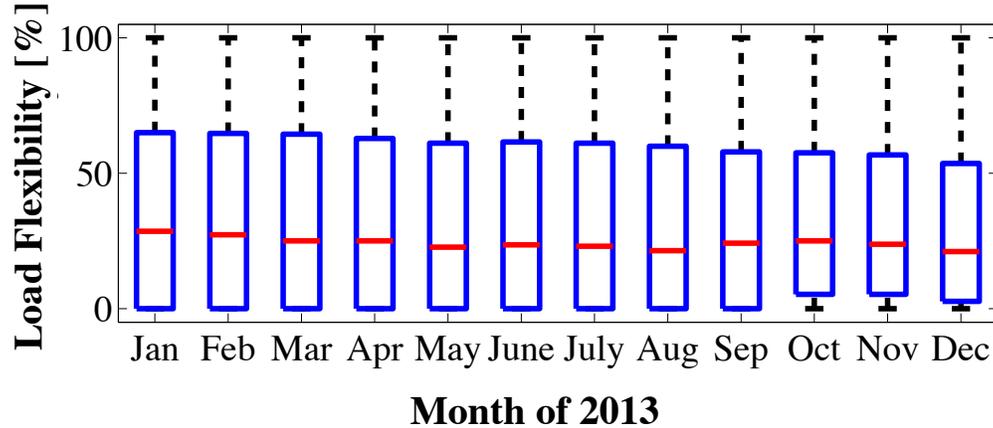


Figure 5.5: The variation in load flexibility

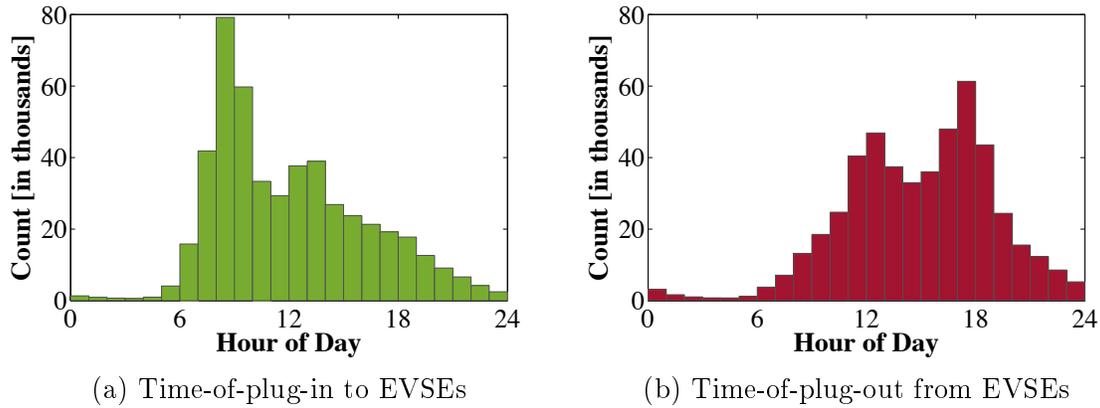


Figure 5.6: Distributions of time-of-plug-in to EVSEs and time-of-plug-out from EVSEs.

5.2 Smart Charging Strategy

In this section, we introduce the proposed smart charging methodology. In particular, we describe the general optimization strategy used to obtain the charging schedules for each charging session.

The goal of the proposed smart charging framework is to reschedule the power time series measured in discrete time slots $[1, \dots, K]$ for any charging session in a population of EVs, $[P_1, P_2, \dots, P_K]$ such that an objective function is optimized. The objective function should capture the desired benefits from a stakeholder's

perspective. While rescheduling the charging, we would like to ensure that the order of the measurements in this time series is preserved. This is because the power that EVSEs draw is dependent on the state of charge (SOC) of the EV that is being charged, and keeping the order of the measurements accounts for this dependency. In addition, we assume that the charging is preemptive; that is, the charging tasks are interruptible without any decrease in the SOC of the EV.

In a typical charging session, an EV starts charging when it is plugged in, and often the charging is complete before the vehicle departs. The smart charging framework proposed in this study is designed to move some of the charging to the slack time slots (i.e. the time slots where the vehicle is plugged in but the charging is completed).

For the purposes of this chapter, we discretize a day into 15-minute intervals. We define the time period for the optimization within a day as the time between the start time slot t_{start} and the end time slot t_{end} . In this period, each charging session i has an arrival time slot denoted by $t_a^{(i)}$ and a departure time slot $t_d^{(i)}$. For each charging session, a column vector including the charging power time series can be created using the power measurements for every time slot in $[t_a^{(i)}, t_d^{(i)}]$. If necessary, the time series is zero-padded to match the size of the optimization time period $[t_{start}, t_{end}]$. Hence, for each EV i , the power time series is given as follows:

$$\mathbf{P}^{(i)} = [P_1^{(i)}, P_2^{(i)}, \dots, P_K^{(i)}]^T \quad (5.3)$$

where K is the total number of time slots in $[t_{start}, t_{end}]$. Next, for each charging session i , we identify $\mathbf{Q}^{(i)}$ whose elements $Q_j^{(i)}$ correspond to the j^{th} non-zero element of $\mathbf{P}^{(i)}$. The goal is to reschedule the time slots $t_j^{(i)}$ in $[t_a^{(i)}, t_d^{(i)}]$ corresponding to $Q_j^{(i)}$ without changing their order. We define $M^{(i)}$ as the total number of non-zero

power measurements in charging session i (i.e. total number of elements in $\mathbf{Q}^{(i)}$).

To capture the precedence and the session duration constraints we proposed above, the following formal constraints are introduced:

$$\left. \begin{array}{l} t_j^{(i)} \geq t_{start} \\ t_j^{(i)} \leq t_{end} \\ t_j^{(i)} \geq t_a^{(i)} \\ t_j^{(i)} \leq t_d^{(i)} \\ t_j^{(i)} < t_{j+1}^{(i)} \end{array} \right\} \begin{array}{l} \forall i \in [1, N], \\ \forall j \in [1, M^{(i)}] \end{array} \quad (5.4)$$

The proposed constraints are constructed using a binary decision matrix to represent charging or non-charging time slots within the optimization duration. In particular, for each element $Q_j^{(i)}$ in $\mathbf{Q}^{(i)}$, we create a binary vector $x^{(i,j)}$ that includes K binary decision variables. Each element in this vector represents a candidate time slot at which $Q_j^{(i)}$ could be positioned. Hence, we define row vectors $x^{(i,j)} \forall i \in [1, N]$ and $\forall j \in [1, M^{(i)}]$. The elements in these vectors are $x_k^{(i,j)} \in \{0, 1\}$ that are defined $\forall k \in [1, K]$.

From these binary vectors $x^{(i,j)}$, we form a binary decision matrix $\mathbf{X}^{(i)}$ for each charging session $i \in [1, N]$. In particular, the individual decision variables $x_k^{(i,j)}$ form the elements of the binary decision matrix $\mathbf{X}^{(i)}$ as follows:

$$\mathbf{X}^{(i)} = \begin{bmatrix} x_1^{(i,1)} & \dots & x_K^{(i,1)} \\ \vdots & \ddots & \vdots \\ x_1^{(i,M^{(i)})} & \dots & x_K^{(i,M^{(i)})} \end{bmatrix} \quad (5.5)$$

Finally, we write the variables in the constraints given in (5.4) using the binary

decision variable as follows:

$$t^{(i)} = \mathbf{X}^{(i)}O, \text{ where } O = \begin{bmatrix} 1 \\ 2 \\ \vdots \\ K \end{bmatrix} \quad (5.6)$$

The aggregate power vector for the VAP $AP^{(d)} = \sum_{i=0}^N(\mathbf{P}^{(i)})$ for the day d is given as follows:

$$AP^{(d)} = \begin{bmatrix} \mathbf{Q}^{(1)} \\ \mathbf{Q}^{(2)} \\ \vdots \\ \mathbf{Q}^{(N)} \end{bmatrix}^T \begin{bmatrix} \mathbf{X}^{(1)} \\ \mathbf{X}^{(2)} \\ \vdots \\ \mathbf{X}^{(N)} \end{bmatrix} \quad (5.7)$$

For each case study proposed in this chapter, we build on the general optimization framework described above, identify the objective functions to capture the benefits from each stakeholder's perspective and introduce additional constraints when necessary. We use the Gurobi optimizer [82] to solve the optimization problems formulated for each case study. Due to the size of the optimization problem for certain VAPs and the number of times the optimization problem is solved to obtain values to estimate benefits for the year of 2013, a proved optimal solution is expected to be hard to reach within a reasonable time frame. For these reasons, we alter the optimality criteria by controlling the relative gap between a feasible integer solution and the general optimal solution. We set this optimality criteria to 5% and allow early termination once a feasible solution is found.

5.3 Charging Infrastructure Owner's Perspective

In the first case study, our goal is to capture and maximize the benefits of smart charging from an EV charging service provider's perspective. Currently, each charging meter is independently owned by the building owner, and the consumption is billed to the building owner as part of the building's monthly bill. However, in our work, we focus only on the load resulting from EV charging, i.e. decoupled from other loads, but aggregated over VAPs formed based on sub-LAPs. This corresponds to the situation in which the charging stations within each VAP are combined and operated under a single owner or an aggregator and the owner is charged according to a time of use (TOU) tariff structure, where shifting load from high cost periods to lower cost periods can offer some benefits to the owner. Although the current VAPs are created based on sub-LAPs, the current scale of the charging infrastructure and the number of charging sessions can easily represent a large parking structure or a campus in the future, where the EV aggregation is behind a single meter and the non-EV load is relatively steady.

5.3.1 Problem Formulation

In a typical TOU rate structure, there are two separate charges forming the monthly bill: the *energy charges* and the *demand charges*. The energy charges are calculated based on the amount of energy consumed over given time periods of the day using the corresponding hourly TOU energy rate. The demand charges are calculated based on the maximum power demand for specific time periods of the day over the course of the billing period. At the end of each billing period, the maximum demand values for the specified periods are multiplied by the demand charge rates and added to the overall energy charge.

In order to model a similar rate structure in the proposed smart charging framework, we define $EC^{(d)}$ as the energy charge for day d of a month with D days (i.e. $d \in [1, \dots, D]$). Then, we define DC_h as the demand charges for each time period h of the day of any month. For example, in PG&E's E-19 TOU rate structure, for winter billing periods, the demand charges are calculated based on 2 time periods *part-peak* (i.e. 8:30AM-12:00PM & 6:00PM-09:30PM) and *off-peak* (i.e. 09:30PM-08:30AM) [6]. Formally, the monthly bill for the owner is therefore given by:

$$f(DC_h, EC^{(d)}) = \sum_{\forall h} DC_h + \sum_{\forall d} EC^{(d)} \quad (5.8)$$

The energy charges $EC^{(d)}$ can easily be incorporated into the proposed daily optimization routine. Defining ER as a column vector reflecting the price of energy for each time slot j , $EC^{(d)}$ for any day d in a billing period is given by:

$$EC^{(d)} = AP^{(d)} ER \quad (5.9)$$

For time period h within day d , a subset of the entire daily aggregate power vector $AP^{(d)}$ is needed and is referred to as $AP_h^{(d)}$.

In order to minimize the cost function given in (5.8), the maximum demand for the daily time periods h must be accurately known beforehand for the entire month. However, in a real life scenario, this is not a valid assumption. To incorporate demand charges into the proposed daily smart charging framework, we therefore propose the following strategy for the owner: for each day d , we define the peak aggregate power values for each period h as $AP_{peak,h}^{(d)}$. Since the historic $AP_{peak,h}$ values for each day in $[1, \dots, d - 1]$ are available to the main scheduler, we can

define the maximum of the historic $AP_{peak,h}$ values until $d - 1$ as follows:

$$AP_{max,h}^{(d-1)} = \max(AP_{peak,h}^{(1)}, \dots, AP_{peak,h}^{(d-1)}) \quad (5.10)$$

Using the above definition, the monthly demand charges can be calculated at the end of the month based on $AP_{max,h}^{(D)}$ and the demand rates DR_h for each period as:

$$DC_h = AP_{max,h}^{(D)} DR_h \quad (5.11)$$

As we move from one day to the next, we try to limit the demand charges based on the maximum daily demands occurred up to the current day. At the beginning of the billing period, we start with no knowledge of the historical peak values, and we keep track of the maximum historical value up to day d . This strategy can be represented by incorporating the maximum value of the peak values $AP_{max,h}^{(d)}$ for time period h and day d as decision variables into the following optimization problem:

$$\underset{\mathbf{X}^{(i)}, AP_{max,h}^{(d)}}{\text{minimize}} \quad AP_{max,h}^{(d)} DR_h + EC^{(d)}$$

subject to (5.4) and the following additional constraints:

$$\left. \begin{array}{l} AP_{max,h}^{(d-1)} \leq AP_{max,h}^{(d)} \\ AP_h^{(d)} \leq AP_{max,h}^{(d)} \end{array} \right\} \forall h \in [1, TP] \quad (5.12)$$

Note that with (5.12), we ensure that the current maximum $AP_{max,h}^{(d)}$ is more than or equal to the maximum historical value $AP_{max,h}^{(d-1)}$ for period h . By definition, this allows for the tracking of the maximum value up to that day. In addition, these maximum values set the day based on which the demand charges will be

Demand Charges	\$/kW	Time Period
Max. Peak Demand Summer	\$19.71253	12:00PM-6:00PM
Max. Part-Peak Demand Summer	\$4.07	8:30AM-12:00PM & 6:00PM-09:30PM
Max. Demand Summer	\$12.56	Any time
Max. Part-Peak Demand Winter	\$0.21	8:30AM-09:30PM
Max. Demand Winter	\$12.56	Any time
Energy Charges	\$/kWh	Time Period
Peak Summer	\$0.16253	12:00PM-6:00PM
Part-Peak Summer	\$0.11156	8:30AM-12:00PM & 6:00PM-09:30PM
Off-Peak Summer	\$0.07818	09:30PM-08:30AM
Part-Peak Winter	\$0.10479	08:30AM-09:30PM
Off-Peak Winter	\$0.08200	09:30PM-08:30AM

Table 5.2: E-19 rate structure [6]

calculated. If none of the current peak values exceeds the historical maximum values, the demand charges for each period h are not set by the current day d .

5.3.2 Case Study

For the purposes of this chapter, we use the demand and energy rates from PG&E's E-19 TOU rate structure [6]. The E-19 rate structure gives the owner the option to manage their electric costs by shifting load from high cost periods to lower cost periods. Detailed information on E-19 is given in Table 5.2. The summer period starts with May 1st and ends October 31st, and the winter period includes the remaining months of the year. This rate is for non-residential customers in PG&E's territory with highest demand exceeding 499 kW for three consecutive months.

To evaluate the benefits of smart charging when the EV aggregation has a single bill calculated on a TOU tariff, we first calculate the current bill under this tariff but without smart charging. Then, we use the proposed optimization strategy to

schedule the loads in a way that minimizes the customer's monthly bills, and we report each monthly bill calculated for each VAP and the contributions from energy and demand charges in the bill.

5.3.3 Results

Figure 5.7 shows the sum of monthly bills calculated in dollars for all of the VAPs. For each month, the left bar shows the current bill, and the right bar shows the optimized bill for the month. It is obvious that the difference between the summer and winter rates impacts the aggregate monthly bill. The increase within the winter and the summer period is due to the increase in the number of charging sessions over the year.

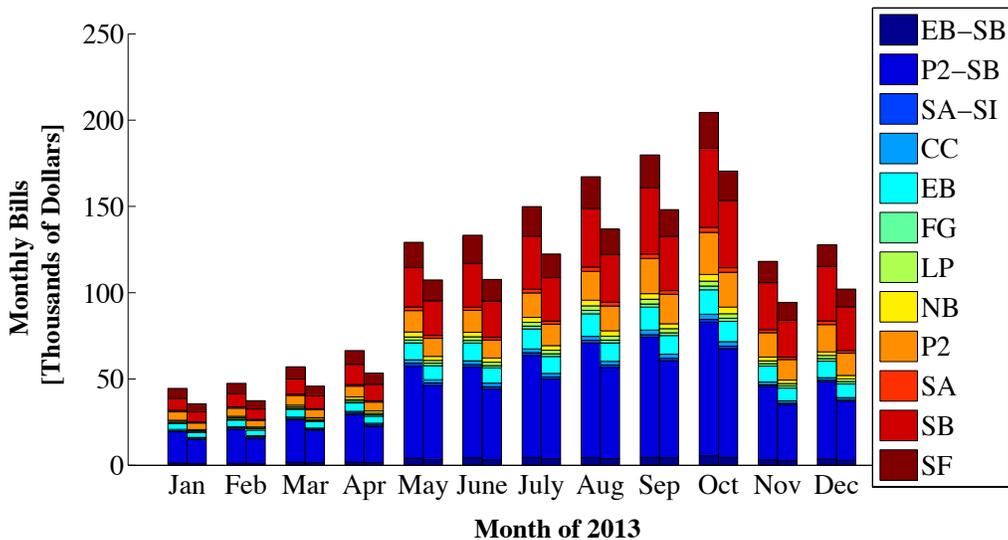
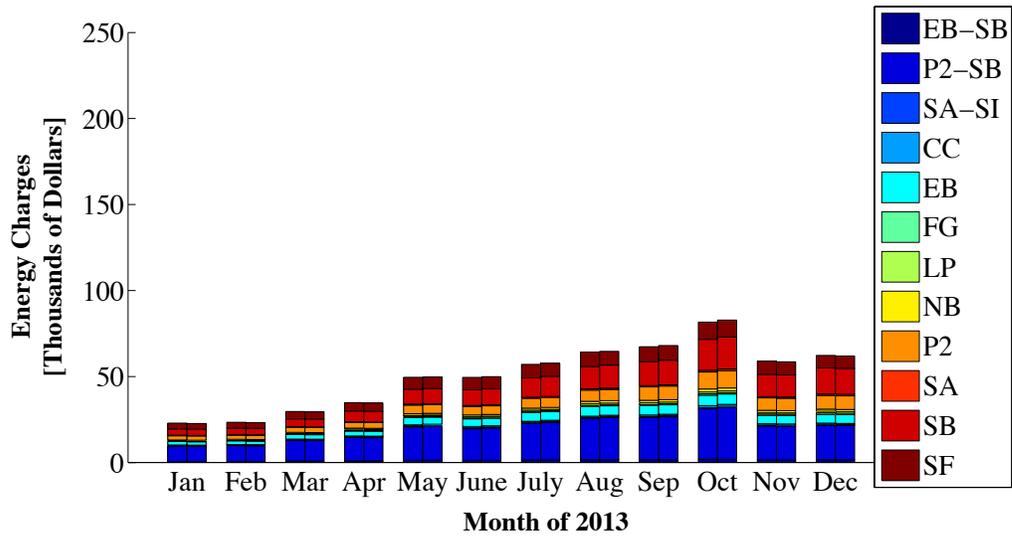
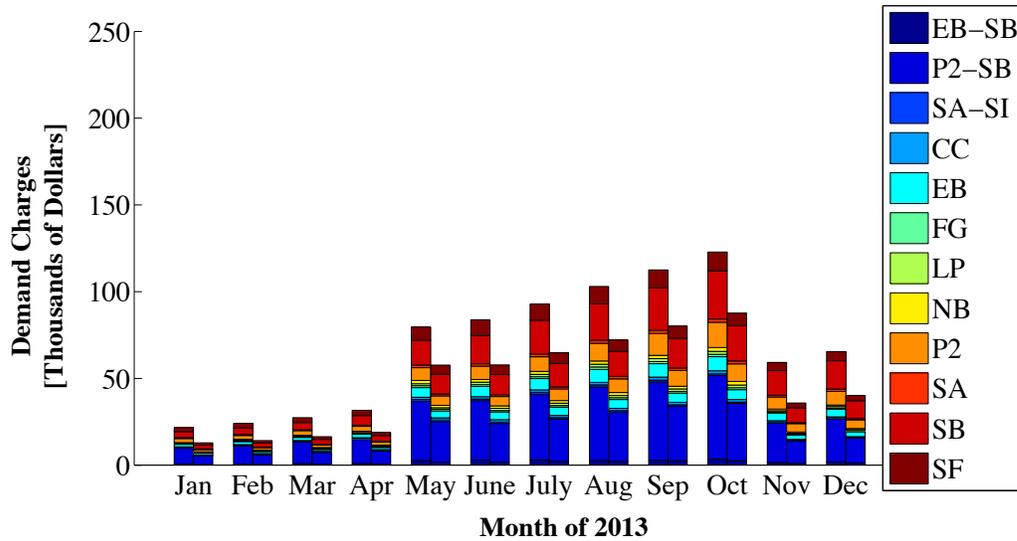


Figure 5.7: Monthly bills calculated with E-19. The left bar for each month shows the current bill, and the right bar shows the optimized bill.

Figures 5.8a and 5.8b show the total energy and demand charges, respectively, over all LAPs. The cumulative energy charges increase slightly for the summer months when using smart charging, whereas there is a significant drop in the demand charges. This suggests that the peak load of the EVs is shifted from the morning



(a) Monthly energy charges calculated with E-19



(b) Monthly demand charges calculated with E-19

Figure 5.8: Decomposition of monthly bills to energy and demand charges. In each figure, the left bar shows the current charges, and the right bar shows the optimized charges for each month.

partial-peak period to the peak-period. This shift is still beneficial because the increase in the energy charges is significantly lower than the decrease in the demand charges.

The cumulative load shapes given in Figure 5.3 and the plug-in and plug-out

VAP	Period	Bill [dollars]		Reduction [dollars /session]	Reduction [%]		
		Current	Optimized		DC	EC	Total
P2-SB	Summer	63001	50395	0.65	20.86%	-0.85%	20.01%
	Winter	29603	22575	0.46	23.41%	0.33%	23.74%
EB-SB	Summer	4588	3788	0.52	16.96%	0.49%	17.45%
	Winter	2092	1724	0.28	17.23%	0.36%	17.59%
SA-SI	Summer	1645	1413	0.36	13.80%	0.30%	14.10%
	Winter	828	752	0.13	9.06%	0.12%	9.18%
CC	Summer	2365	2178	0.24	7.34%	0.57%	7.91%
	Winter	1037	896	0.22	13.31%	0.29%	13.60%
EB	Summer	12033	10003	0.41	16.44%	0.43%	16.87%
	Winter	5874	4868	0.26	16.66%	0.47%	17.13%
FG	Summer	1803	1568	0.33	11.98%	1.05%	13.03%
	Winter	920	807	0.18	11.82%	0.46%	12.28%
LP	Summer	2370	2135	0.29	9.37%	0.55%	9.92%
	Winter	1141	1002	0.20	11.88%	0.30%	12.18%
NB	Summer	3136	2865	0.23	8.16%	0.49%	8.64%
	Winter	1391	1271	0.13	8.48%	0.22%	8.63%
P2	Summer	16795	14171	0.48	16.13%	-0.51%	15.62%
	Winter	8567	7010	0.34	17.98%	0.20%	18.17%
SA	Summer	2313	1991	0.45	13.88%	0.04%	13.92%
	Winter	1215	914	0.52	24.76%	0.01%	24.77%
SB	Summer	32911	27439	0.53	17.72%	-1.09%	16.63%
	Winter	15645	12602	0.37	19.34%	0.11%	19.45%
SF	Summer	17679	14224	0.51	18.07%	1.47%	19.54%
	Winter	8591	7046	0.28	17.10%	0.88%	17.98%

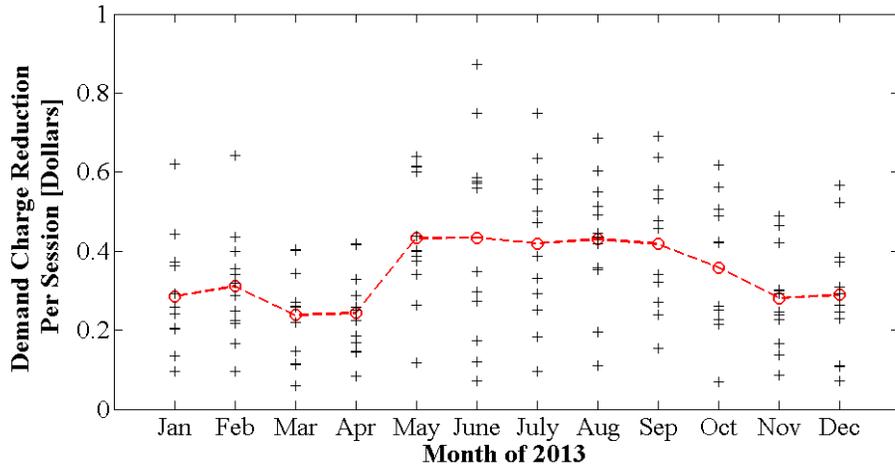
Table 5.3: Average results based on summer and winter month rates in E-19

time histograms given in Figures 5.6a and 5.6b support these results. These figures suggest that energy charges increase because a large portion of the EV charging sessions end (i.e. the charger is unplugged) before the system peak period ends. Thus, when coupled with the higher part-peak demand rates, the optimization converges to a result in which the load is shifted from the EV load peak period (9AM-11AM) to the system peak period (12PM-6PM).

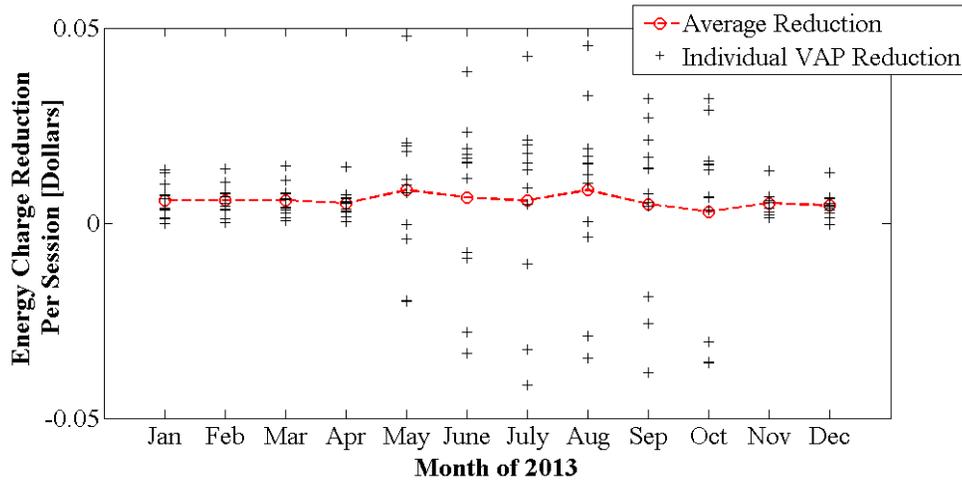
The results given in Table 5.3 provide further insight into the results depicted in Figures 5.7, 5.8a and 5.8b. Specifically, we reflect on the average monthly bill

before and after optimization for winter and summer months. Then, we report on average bill reduction per session during these periods. The values range between 0.13 and 0.65 dollars among all VAPs. Overall, we find that the rate structure in the summer periods yields to more reductions per session than the rates in winter months, with the exception of the Sacramento Valley (SA) VAP. We also report on the total percent bill reduction and we break down this percentage into contributions from demand charges and energy charges. We observe that average percent bill reductions range between 8.63% and 24.77%. Even though the average reduction per session values are mostly higher during summer months, the relative cost reduction in monthly bills for individual VAPs varies less. This is due to high overall costs in the summer months.

To further understand the relationship between the monthly bill reductions and the different rate structure between summer months and winter months, we include Figures 5.9a and 5.9b. Specifically, we depict the charging-session-normalized demand charge and energy charge reduction values for each VAP for all months of 2013 in Figures 5.9a and 5.9b, respectively. The dashed lines represent the average charging-session-normalized reduction value for both plots. When winter and summer months are investigated separately, we observe that the charging-session-normalized reduction values remain similar to each other. Per charging session, the expected economic benefit from the demand charge reduction is much higher than the benefit from energy charge reduction. However, differences exist between the reduction results obtained for winter and summer months. In Figure 5.9a, we observe that the average demand charge reduction per session is higher in summer months than the winter months. We believe that this is due to higher demand charges in the summer months. In Figure 5.9b, we observe that the energy charge reductions are



(a) Demand Charge Reduction in Dollars per Session

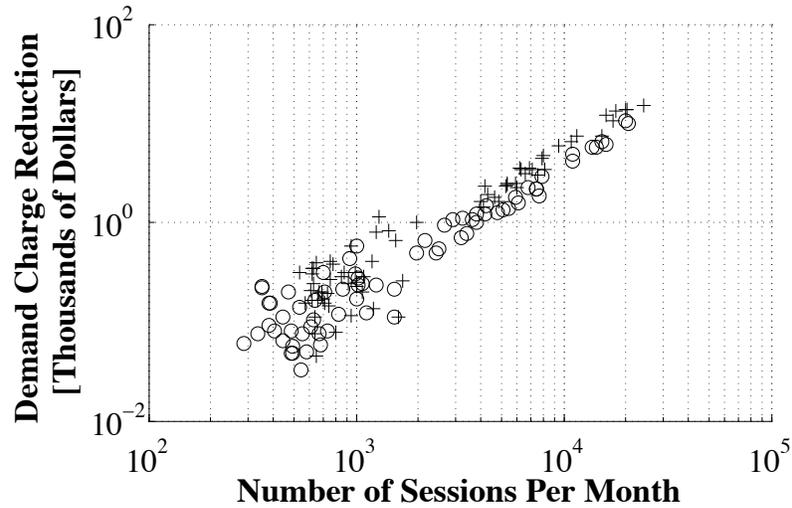


(b) Energy Charge Reduction in Dollars per Session

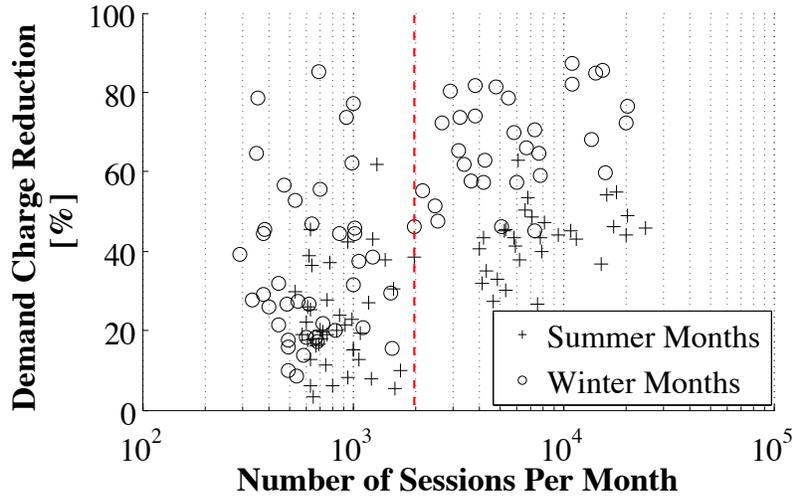
Figure 5.9: Charging-session-normalized Demand and Energy Charge Reductions

mostly non-negative for winter months. However, for summer months we observe negative energy charge reduction values per charging session. This is in line with the results depicted in Figure 5.8a. An increase in the energy charge is observed for the summer periods due to shifting energy from part-peak periods to the peak periods.

Figures 5.10a and 5.10b depict the relationship between the reduction in demand



(a) Demand Charge Reduction in Dollars by Session Size



(b) Demand Charge Reduction Percentage by Session Size

Figure 5.10: Demand Charge Reduction by Session Size

charges and the number of charging sessions in each VAP per month. Specifically, in Figure 5.10a, we examine the decrease in demand charges in dollars. We observe a linear trend: as the number of sessions per month rises, the reduction in demand charges increases linearly. Given the current load flexibility and plug-in and plug-out times, this is expected because most of the EVs contribute to the peak load of the EV aggregation. In Figure 5.10b, we look at the percent reduction in demand

charges. For up to 2000 charging sessions per month (indicated by a red dashed line in Figure 5.10b), there is no clear separation between the winter and summer months and, for a given number of sessions, the demand charge reduction values vary. Beyond this point, we can see a clear separation between the winter and summer months, and the demand charge reduction values show less variance.

The relative decrease in the summer months is less than the relative decrease in the winter months. We believe that this is due to the time of the peak EV load, the plug-in and plug-out patterns of the EVs and the corresponding rate structure. In particular, the peak EV load coincides with the part-peak rate period, and most of the EVs depart before the system peak period (12PM-6PM) is over. The system peak period has a separate and higher demand rate in the summer months (detailed in Table 5.2). This limits the smart charging framework's ability to move the EV loads from part-peak period to system peak period. The winter rates we use in this study do not include a separate demand rate for the system peak period; rather, the part-peak period extends from 8:30AM-09:30PM. This makes it possible to manage the EV peak load in a more effective way.

5.4 Distribution System Operator's Perspective

In the second case study, we evaluate the potential benefits that smart charging can offer to distribution system operators (DSOs). The motivation behind this case study is to investigate the potential of each charging session to decrease its contribution to the peak system demand via smart charging. We first quantify the percentage of peak load shed during the system peak load period (12-6PM). We then quantify the amount of energy that is shifted outside the peak period by the EV load aggregation for each month of 2013. Finally, we report on the amount of

energy that can be expected to be moved outside of the system peak period per charging session.

5.4.1 Problem Formulation

To realize peak shaving, we propose to develop a two-stage optimization. In the first stage, we minimize a bound on the aggregate power consumed by the EVSEs within a VAP during the pre-defined peak period (12-6PM) only. We simply refer to the pre-defined peak period as pp , and to simplify the notation introduced earlier, we refer to the aggregate power vector within the peak period as $AP_{pp}^{(d)}$. To implement the initial stage optimization, we define $AP_{bound,pp}^{(d)}$ as a decision variable to represent the proposed bound on the $AP_{pp}^{(d)}$. Then, in the second stage, using the optimal bound as a constraint, we minimize the total energy consumed within the peak period. This implicitly ensures that the energy bill for the customer is decreased or unchanged based on a typical TOU tariff.

The first part of the optimization can be written as:

$$\underset{\mathbf{X}^{(i)}, AP_{bound,pp}^{(d)}}{\text{minimize}} \quad AP_{bound,pp}^{(d)}$$

subject to (5.4) and the following additional constraints:

$$AP_{pp}^{(d)} \leq AP_{bound,pp}^{(d)} \quad (5.13)$$

Then, using the optimal $AP_{bound,pp}^{(d)}$ values obtained in the first stage $AP_{bound,pp}^{*(d)}$, we can form the second stage as follows:

$$\underset{\mathbf{X}^{(i)}}{\text{minimize}} \quad \sum_{\forall k \subseteq pp} AP_k^{(d)}$$

subject to (5.4) and the following additional constraints:

$$AP_{pp}^{(d)} \leq AP_{bound,pp}^{*(d)} \quad (5.14)$$

5.4.2 Case Study

The motivation behind our second case study is to evaluate the potential of EV aggregations to decrease their contribution to the system peak load via smart charging. As the plug-in and plug-out time histograms given in Figures 5.6a and 5.6b suggest, the amount of energy that can be moved outside of the peak period is expected to be low, mostly because most non-residential EV sessions end before the system peak period is over. However, there is potential in using smart charging and exploiting the inherent flexibility in each charging session to decrease the contribution of EVs to the system peak load.

To demonstrate and quantify this potential, we calculated optimal schedules for each VAP-level aggregation using the optimization strategy described in the above section, and obtained percentage of peak shed values and the total energy moved outside of the peak period for every day in each month of 2013.

5.4.3 Results

Figure 5.11 shows the box plots created using daily peak shed values for each month of 2013. The percentage of peak shed for each day d is defined as:

$$\%peakshed^{(d)} = \frac{AP_{bound,pp}^{*(d)}}{\max(AP_{pp}^{(d)})} \quad (5.15)$$

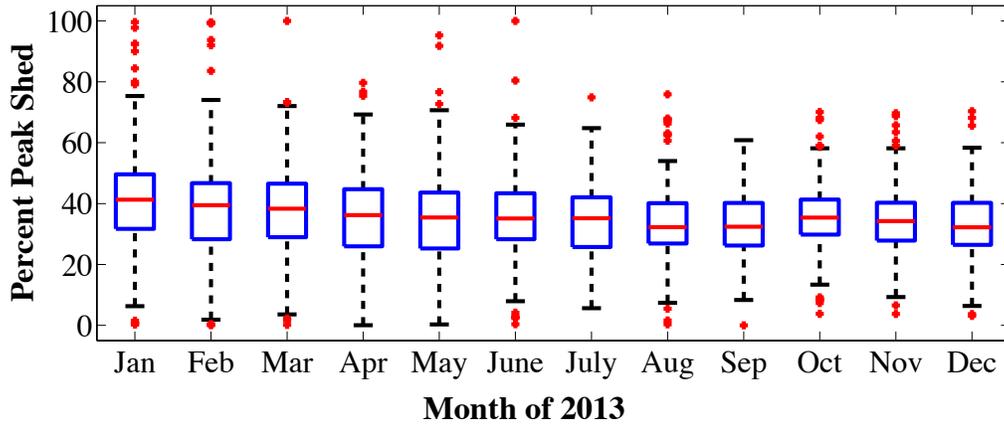


Figure 5.11: Distribution of percent peak shed for all the VAPs

The red lines denote the median value of the distribution, the box boundaries are the 25th and 75th percentiles and the whiskers denote the 1st and 99th percentiles, assuming the distributions per each month are normal. The outliers outside the whiskers' boundaries are marked with points. As expected, the smart charging significantly reduces the peak EV load during the system peak period. The median values for all of the months range between 30 and 42%. A decrease in the peak shaving potential and a slight decrease in the variation of the distributions over the course of 12 months are also apparent in Figure 5.11. This can be explained by the increase in the number of charging sessions per EVSE and the related decrease in the variation of available flexibility.

Figure 5.12 depicts the distribution of the average energy moved outside of the peak period per charging session for all of the VAPs estimated every day of the month. The median value over 2013 is approximately 0.25kWh per charging session, which corresponds to $\sim 2.8\%$ of the average energy put during each charging session.

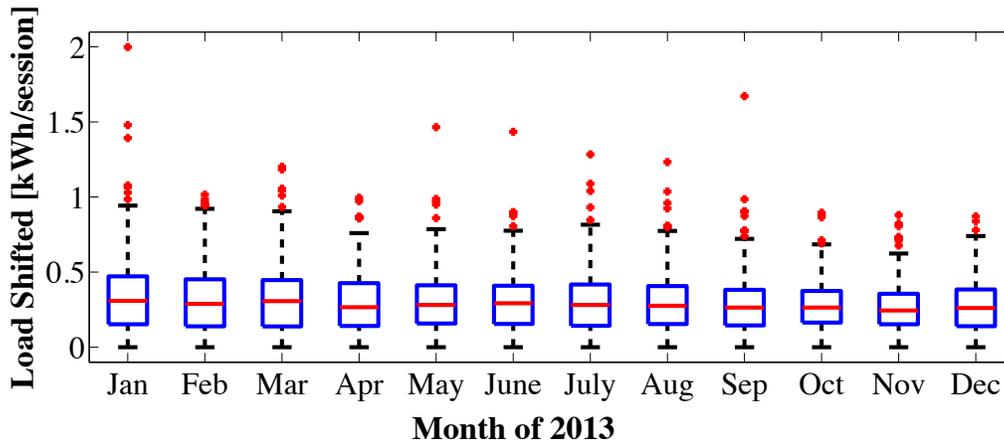


Figure 5.12: Total energy moved outside of system peak period

5.5 Chapter Conclusions

In this chapter, we quantify the potential benefits of smart charging to different stakeholders using data collected from over 2000 non-residential electric vehicle supply equipment (EVSEs) located throughout 190 zip code regions in Northern California. We created virtual aggregation points (VAP) in which the aggregate power consumption of a selected population of EVSEs is assumed to be managed via individual charging control at each EVSE. We developed and used a smart charging framework to estimate the benefits of EV smart charging to different stakeholders: a single owner/an aggregator of behind-the-meter EVSEs (i.e. aggregators) and distribution system operators.

In our first case study, we investigated the potential benefits of behind-the-meter EV aggregations. The aggregate load is re-scheduled using a TOU rate structure. Our results suggest that up to 24.8% decrease in the aggregate monthly bill per VAP is possible. In all months, this reduction is due to a corresponding decrease in demand charges in the monthly bill: we observed that decreases in energy charges are contributing by up to 1.5% to the overall decrease, whereas the demand charges

contribute up to 24.7%.

In our second case study, we used the EV aggregations to decrease their contribution to the system-level peak load. We have observed median peak shed values around 30%-42% for each month. In addition, we have quantified the amount of energy that can be shifted outside the peak period per charging session over the course of 2013, and found the median value to be approximately 0.25kWh/session ($\sim 2.8\%$ of the average energy put in each session).

Chapter 6

Conclusions

6.1 Key findings

In this dissertation, we focused on the use of data obtained from sensors to improve demand-side management techniques that support the reliable operation of the power grid. In particular, we envisioned developing demand response strategies that make use of the inherent flexibility of end use loads while ensuring that these strategies are non-disruptive to the users. We made use of data obtained from smart devices that can monitor and control end use loads, specifically thermostatically controlled loads (TCLs) and electric vehicles (EVs).

In Chapter 3, we tested the following hypothesis: **a Markov Decision Processes (MDP) based model of a heterogeneous TCL population decreases the plant-model mismatch in comparison to a state-of-the-art linear-time-invariant model, when disturbances to individual TCL operations are considered.** First, we developed an individual TCL modeling strategy which allowed us to include disturbances to their individual dynamics when modeled. Building on this model, we created an aggregate model using a bottom-up approach inspired by

an MDP-based strategy. We quantified the plant model mismatch of the proposed strategy, and compared it with a benchmark linear-time-invariant model from the literature using a realistic simulation of household refrigeration units (HRUs).

Our results suggest that if the plant is simulated or sampled close to real time, the assumptions made in the benchmark linear-time-invariant model holds true. However, neither information on the thermal resistance and capacitance values of the loads nor historical power consumption data from these loads may be available. In addition, there may be limits or delays on the communication infrastructure during training, hence the real-time state observations of the population may be unavailable. The proposed MDP-based strategy brings extra value in such cases.

Although the timing requirements of different ISOs or RTOs participating in ancillary services differ by 2 to 10 seconds [19, 50, 51], the requirement decreases to 5 to 10 minutes for services such as load following [19]. Depending on the service to be provided (i.e. the timing requirement) and the modeling accuracy needed, the proposed MDP-based strategy can be used to better capture the dynamics of the aggregate HRU population.

Finally, the proposed MDP based strategy creates no computational burden compared to the benchmark LTI system. Our results suggest that the benchmark LTI system acts as a boundary case for the proposed MDP-based system, and the additional degree of freedom brought by the MDP-based system brings extra value in specific cases when the simulation time steps are higher and non-linearities become more evident.

In Chapter 4, **we quantified the accuracy of information on the state of an aggregate TCL population that can be inferred from aggregate power measurements obtained at a neighborhood level using the proposed**

MDP-based modeling strategy.

We first developed a moving horizon mean squared error state estimator with constraints to estimate the states of a heterogeneous TCL population from aggregate power measurements of a TCL population. We analyzed the performance of the proposed MHSE under different time horizon T values and compared the performance of this estimator with a generic Kalman filtering approach based on [3] under different switching conditions imposed by the forcing parameter f . The results show that the MHSE performs better in scenarios where the expected switching frequency is higher than 25%. We believe that this increase is due to the nature of the model used for the MHSE, specifically due to the additional flexibility introduced in the model through changing the way control actions are included in the system dynamics. An increase in the switching frequency is often expected when the controller is trying to keep a constant aggregated power consumption over certain period of time, and when there are abrupt changes in the reference power followed by the TCL population.

Next, we relaxed the assumption regarding the availability of the TCL population's aggregate power measurements. Specifically, we introduced the contextually supervised state estimator (CSSE) to infer the states of a heterogeneous TCL population from power measurements obtained at a neighborhood level distribution feeder. As a proof of concept, we showed that the CSSE outperforms a simple prediction strategy in cases where the CSSE has information on the thermal properties of the individual TCLs in the load population. Our results also suggest that in such cases, the increasing proportion of the non-TCL loads present in the virtual aggregation point decreases the performance of the CSSE, however we still see significant improvements in the CSSE performance in comparison to the simple prediction sce-

nario for cases where the TCL load is approximately 20% of the non-TCL load in the aggregate power signal. Finally, we investigated a case where the CSSE had no information on the thermal properties of the individual TCLs in the load population and it uses the parameters available in the literature to perform system identification. The results of this investigation suggests that the CSSE's performance is highly dependent on the availability of a reliable model that uses information on the thermal properties of individual TCLs in the population.

In Chapter 5, we quantified the benefits of managed EV charging to distribution system operators and EV charging infrastructure owners, by leveraging an extensive dataset of more than 500,000 commercial EV charging sessions.

Using data collected from over 2000 non-residential electric vehicle supply equipments (EVSEs) located in Northern California for the year of 2013, we developed a smart charging framework to identify the benefits of non-residential EV charging to the load aggregators and the distribution grid. With this extensive dataset, we improved upon past studies focusing on the benefits of smart EV charging by relaxing the assumptions made in those studies regarding: (*i*) driving patterns, driver behavior and driver types; (*ii*) the scalability of a limited number of simulated vehicles to represent different load aggregation points in the power system with different customer characteristics; and (*iii*) the charging profile of EVs.

First, we studied the benefits of EV aggregations behind-the-meter, where we used a time-of-use rate structure to understand the benefits to the owner when EV aggregations shift load from high cost periods to lower cost periods. For the year of 2013, we showed that a reduction of up to 24.8% in the monthly bill is possible. Then, following a similar aggregation strategy, we showed that EV ag-

gregations decrease their contribution to the system peak load by approximately 40% when charging is controlled within arrival and departure times. Our results also showed that it is possible to shift approximately 0.25kWh ($\sim 2.8\%$) of energy per non-residential EV charging session from peak periods (12PM-6PM) to off-peak periods (after 6PM) in Northern California for the year of 2013.

6.2 Discussion

In every chapter of this dissertation, we carefully define the user's expectations of the loads. In Chapters 3 and 4, for example, we assume that the thermostatic deadband represents the *comfort region* of the user, and as long as it is not violated, we consider the user *comfortable*. Similarly, in Chapter 5, we assume that the arrival and departure times of the EVs and the amount of energy put in each charging session remain the same. Therefore, for both TCLs and EVs, the results claimed in this dissertation inherently assume that these assumptions reflect realistic representations of end user's comfort and expectations of these loads.

In Chapter 3, we assume that the TCLs are aware of their internal states, that they can receive a broadcasted control signal, and that they have an interface through which upstream communication is possible. Today, most loads lack the necessary communication capabilities, and their control systems are proprietary. Further development and acceptance of the Internet of Things (IoT) concept can lead to a solution to these challenges; however privacy remains a big concern. End users define privacy in different ways, the trade off between the availability of data, the privacy of end users and the accuracy of availability of DR services must still be investigated.

In Chapter 4, we investigated a case where we limited the upstream information exchange between the TCLs and the main controller to protect privacy and reduce the cost of implementation. Our results showed that if the thermal parameters of a TCL population are known, the current state of the load population can be inferred from an aggregate power measurement of the TCLs or a neighborhood level aggregation point in which non-TCL loads are grouped with the TCLs. We observed great sensitivity to the load model, which is derived based on the thermal properties of the TCLs that are being controlled. Hence, based on our preliminary results, the one-time information transfer of the thermal properties of individual loads remains key in eliminating the subsequent real-time information transfer of state information.

In Chapter 4, we used *random* actions sent to the main controller to perturb the aggregate population of loads, which led to a better state estimation performance as discussed in Section 4.9. However, the following question remains unanswered: *What is the best action/ set of actions (i.e. a policy) that results in a guaranteed state estimation performance?* One simple suggestion could be to broadcast a control signal that turns all the appliances OFF or ON. This would shrink the state space of the true state to half of its original size.

In addition, as stated in Section 4.8, our analysis focused on the performance of the proposed state estimation strategy, rather than on the performance of the overall source separation between the different types of loads modeled. We think that with the current time horizon values for the CSSE strategy, achieving an acceptable source separation accuracy for non-TCL loads might be challenging. For approaches that target load identification from smart meter data, different TCL modeling techniques that are based on exogenous inputs are presented in [52].

In Section 4.8, we provided a proof-of-concept state estimation strategy using the contextually supervised source separation strategy. However, within the scope of Chapter 4, we did not do an extensive tuning of the weighting parameters α, β and γ . A more systematic tuning and further investigation of the closed loop performance of these state estimation strategies are necessary before implementation.

In Chapter 5, we assumed that the managed charging strategy has perfect information over the EVs' arrival and departure times. Thus, the quantified benefits serve as an upper bound for stochastic managed charging approaches, which include more realistic assumptions on the availability of such information. In a real life scenario, this information could either be requested from the driver upon arrival (with certain confidence), or it can be predicted based on historical arrival and departure patterns of the driver.

6.3 Broader Impact

Chapters 2, 3 and 4 have important implications for the role of TCLs in the power grid. Specifically, we investigated the necessary sensing points and communication requirements for different types of modeling strategies. In Chapter 2, we developed a DR potential estimation strategy from a large aggregation of residential HVAC loads. This study provided crucial insights into the capability of TCLs to provide large-scale load shedding with setpoint adjustments using an open loop control scenario. To the best of our knowledge, this is the first work that uses energy consumption and temperature data obtained at the device level to estimate models that can then be used to compute DR potential at this scale (over 4,000 households). This study helped us understand the duration of a load shed event initiated by a setpoint adjustment and the likelihood of failure to provide the expected load shed

event in time. It also reinforced the need to use aggregate TCL models and state estimation techniques to address computational complexity and privacy concerns about individual load models. We investigated those issues further in Chapters 3 and 4.

In Chapter 3, we focused on aggregate TCL models. Using data collected from appliances, we investigated strategies to better capture the dynamics of the population with the aggregate models when disturbances to individual TCLs are present. Our results have direct implications for the implementation readiness of state-of-the-art TCL control strategies. We showed that our proposed model generalizes over DR services with different time scales (i.e. load following and regulation) and introduces no additional computational burden. We believe that a scheduling entity that participates in fast DR services on different time scales can benefit from the proposed model.

In Chapter 4, we focused on the communication requirements between the end use loads and a centralized controller when a population of TCLs are controlled to provide DR services. We believe that end users will prefer a strategy in which they do not have to share real-time information with the centralized controller. Using measurements obtained at a neighborhood level aggregation point, we developed state estimation techniques to replace the need for upstream information exchange between individual end use loads and the controller. Understanding these communication requirements makes it possible to identify communication architectures that are well-suited for the control of TCLs for DR services. This chapter also has important policy implications for appliance standards and the level of information exchange expected from appliances in the future. The state estimation techniques developed in this chapter can help satisfy the current telemetry and metering re-

quirements in various independent system operator (ISOs) regions at a lower cost than the current strategies, which could boost the adoption of fast DR strategies capable of providing ancillary services. For example, in Texas, the scheduling entity is required to send a telemetry to ERCOT every two seconds [50]. If an aggregation of behind-the-meter loads are participating in fast DR, this telemetry must include the sum of the individual loads under control. The proposed state estimation techniques can help estimate this aggregate value without sub-metering.

In Chapter 5, we quantified the benefits of smart EV charging to various stakeholders in the power grid. To the best of our knowledge, this is the first study that leverages such an extensive EV charging dataset (over 580,000 charging sessions). Managing the EV charging loads makes it possible to provide load-shedding benefits to the distribution system operators and economic benefits to the charging infrastructure owners without any hindrance to end users. We quantified the peak load reduction capacity within the current system peak period for distribution system operators. We also quantified the economic benefits for charging infrastructure owners using a time-of-use rate structure. Our results have direct implications on the understanding of the effectiveness of smart charging under current DR programs and time of use rating structures. We also reveal the relationship between the EV infrastructure availability, EV load flexibility and the benefits to the stakeholders, which provides a clear picture of managed EV charging for future adjustments and improvements in DR programs. Specifically, our results can serve charging infrastructure owners who are deciding whether to invest in additional charging stations, as the results present the economic benefits provided by different numbers of charging stations.

6.4 Future Research Topics

The work presented in this dissertation serves as a starting point for different avenues of research. These include:

- Leveraging the aggregate TCL modeling strategy and the CSSE proposed in this dissertation to investigate closed loop control strategies for controlling TCLs.
- Developing of time-dependent aggregate TCL models to explicitly capture time varying exogenous inputs and the disturbances to the load population.
- Conducting pilot studies to demonstrate the capabilities of the proposed state estimation and modeling techniques.
- Coupling cellular data (i.e. call detail records) with the EV charging data to understand the dynamics of EV adoption and range anxiety.
- Leveraging call detail records, spatially referenced census data and the data used in Chapter 5 to understanding EV growth and penetration in California.
- Analyzing and comparing the managed EV charging benefits for non-residential locations with different types of commercial activity (i.e. retail store, office building, public parking, etc.).
- Using the dataset presented in Chapter 5 to investigate stochastic managed EV charging strategies that leverage predicted driver arrival and departure times.

Bibliography

- [1] S. Koch, J. L. Mathieu, and D. S. Callaway, “Modelling and control of aggregated heterogenous thermostatically controlled loads for ancillary services,” *Power Systems Computation Conference*, Aug. 2011.
- [2] E. Kara, Z. Kolter, M. Berges, B. Krogh, G. Hug, and T. Yuksel, “A moving horizon state estimator in the control of thermostatically controlled loads for demand response,” in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, Oct 2013, pp. 253–258.
- [3] J. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *Power Systems, IEEE Transactions on*, vol. 28, no. 1, pp. 430–440, 2013.
- [4] J. Zimmermann, M. Evans, J. Griggs, N. King, H. L., P. Roberts, and C. Evans, “Household electricity survey: A study of domestic electrical product usage,” Intertek, Tech. Rep. R66141, 2012.
- [5] J. L. Mathieu, “Modeling, analysis, and control of demand response resources,” 2013.

- [6] Pacific Gas and Electric Company, “Electric Schedule E-19:Medium General Demand-Metered TOU Service,” 2010. [Online]. Available: http://www.pge.com/tariffs/tm2/pdf/ELEC_SCHEDS_E-19.pdf
- [7] Y. G. Rebours, D. S. Kirschen, M. Trotignon, and S. Rossignol, “A survey of frequency and voltage control ancillary services-Part I: Technical features,” *IEEE Trans. Power Sys.*, vol. 22, pp. 350–357, 2007.
- [8] M. Ilic and Q. Liu, “Toward a framework for sustainable Cyber-Physical energy systems: Frequency regulation revisited,” in *World Congress*, vol. 18, 2011.
- [9] J. Kassakian and R. Schmalensee, “The future of the electric grid: An interdisciplinary mit study,” Technical report, Massachusetts Institute of Technology, Tech. Rep., 2011.
- [10] C. Gellings, M. Samotyj, and B. Howe, “The future’s smart delivery system [electric power supply],” *Power and Energy Magazine, IEEE*, vol. 2, no. 5, pp. 40 – 48, sept.-oct. 2004.
- [11] A. H. R. Faruqi, G. S.S., J. Bode, P. Mangasarian, I. Rohmund, G. Wikler, and S. Glosh, D. Yoshida, “A national assesment of demand response potential,” Federal Energy Regulatory Commission (FERC), Tech. Rep., June 2009. [Online]. Available: www.ferc.gov/legal/staff-reports/06-09-demand-response.pdf
- [12] J. H. Eto, “Demand response spinning reserve demonstration–phase 2 findings from the summer of 2008,” *Lawrence Berkeley National Laboratory*, 2010.
- [13] D. S. Callaway and I. A. Hiskens, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan. 2011.

- [14] S. Kiliccote, M. A. Piette, and G. Ghatikar, "Smart buildings and demand response," *AIP Conference Proceedings*, vol. 1401, no. 1, pp. 328–338, Nov. 2011. [Online]. Available: http://proceedings.aip.org/resource/2/apcpcs/1401/1/328_1
- [15] R. Liu, L. Dow, and E. Liu, "A survey of PEV impacts on electric utilities," in *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*. IEEE, 2011, pp. 1–8.
- [16] D. Wu, D. Aliprantis, and K. Gkritza, "Electric energy and power consumption by light-duty plug-in electric vehicles," *Power Systems, IEEE Transactions on*, vol. 26, no. 2, pp. 738–746, 2011.
- [17] C. Weiller, "Plug-in hybrid electric vehicle impacts on hourly electricity demand in the united states," *Energy Policy*, vol. 39, no. 6, pp. 3766 – 3778, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0301421511002886>
- [18] S. Kiliccote, "Findings from seven years of field performance data for automated demand response in commercial buildings," *Lawrence Berkeley National Laboratory*, 2010.
- [19] S. Kiliccote, S. Lanzisera, A. Liao, O. Schetrit, and M. A. Piette, "Fast dr: Controlling small loads over the internet," *Forthcoming Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings*, 2014.
- [20] M. A. Piette, J. Granderson, M. Wetter, and S. Kiliccote, "Responsive and intelligent building information and control for low-energy and optimized grid integration," in *Presented at the*, 2012.

- [21] H. Farhangi, "The path of the smart grid," *Power and Energy Magazine, IEEE*, vol. 8, no. 1, pp. 18–28, 2010.
- [22] W. Kempton and J. Tomić, "Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy," *Journal of Power Sources*, vol. 144, no. 1, pp. 280–294, 2005.
- [23] M. A. Piette, D. Watson, N. Motegi, and S. Kiliccote, "Automated critical peak pricing field tests: 2006 pilot program description and results," *Lawrence Berkeley National Laboratory*, 2007.
- [24] J. L. Mathieu, T. Haring, and G. Andersson, "Harnessing residential loads for demand response: Engineering and economic considerations," University of Florida, Gainesville, Dec. 2012.
- [25] U. Department of Energy, "2011 Buildings Energy Data Book," 2012. [Online]. Available: http://buildingsdatabook.eren.doe.gov/docs/DataBooks/2011_BEDB.pdf
- [26] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, May 2009.
- [27] J. H. Williams, A. DeBenedictis, R. Ghanadan, A. Mahone, J. Moore, W. R. Morrow, S. Price, and M. S. Torn, "The technology path to deep greenhouse gas emissions cuts by 2050: the pivotal role of electricity," *science*, vol. 335, no. 6064, pp. 53–59, 2012.

- [28] T. A. Becker, I. Sidhu, and B. Tenderich, “Electric vehicles in the united states: A new model with forecasts to 2030,” *Center for Entrepreneurship and Technology, University of California, Berkeley*, no. 2009.1, 2009.
- [29] D. Bank, “Electric cars: plugged in,” *Batteries must be included. Frankfurt (Global Markets Research, 9.06. 2008)*, 2008.
- [30] E. C. Kara, M. D. Tabone, J. S. Macdonald, D. S. Callaway, and S. Kiliccote, “Quantifying flexibility of thermostatically controlled loads for demand response:a data-driven approach,” in *Proceedings of the 1st International Conference on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ser. BuildSys '14. Memphis, TN, USA: ACM, 2014.
- [31] E. C. Kara, M. Berges, B. Krogh, and S. Kar, “Using smart devices for system-level management and control in the smart grid: A reinforcement learning framework,” *IEEE SmartGridComm'12*, Nov. 2012.
- [32] E. C. Kara, M. Bergés, and G. Hug, “Modeling thermostatically controlled loads to engage households in the smart grid: Lessons learned from residential refrigeration units,” *Proceedings of the 2013 ASCE the International Conference on Computing in Civil and Building Engineering*, 2014.
- [33] E. C. Kara, M. Berges, and G. Hug, “Impact of disturbances on modeling of thermostatically controlled loads for demand response,” *Smart Grid Special Issue: Cyber Physical Systems and Security for Smart Grid, IEEE Transactions on*, September 2014 (Under Review).
- [34] E. C. Kara, J. S. Macdonald, D. Black, M. E. Berges, G. Hug, and S. Kiliccote, “Estimating the benefits of smart charging an aggregation of electric vehicles in

- commercial buildings: A data-driven approach,” *Applied Energy*, 2014 (Under Review).
- [35] N. Lu, D. Chassin, and S. Widergren, “Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model,” *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 725–733, May 2005.
- [36] J. Short, D. Infield, and L. Freris, “Stabilization of grid frequency through dynamic demand control,” *Power Systems, IEEE Transactions on*, vol. 22, no. 3, pp. 1284–1293, Aug. 2007.
- [37] S. Bashash and H. Fathy, “Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads,” in *American Control Conference (ACC), 2011*, 2011, pp. 4546–4553.
- [38] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, “Frequency regulation from flexible loads: Potential, economics, and implementation,” 2013.
- [39] K. Turitsyn, S. Backhaus, M. Ananyev, and M. Chertkov, “Smart finite state devices: A modeling framework for demand response technologies,” Mar. 2011. [Online]. Available: <http://arxiv.org/abs/1103.2750>
- [40] J. Taneja, K. Lutz, and D. Culler, “The impact of flexible loads in increasingly renewable grids,” in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*. IEEE, 2013, pp. 265–270.
- [41] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld, “Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations,” *Proceedings of the ACEEE Summer Study on Buildings, Pacific Grove, CA*, 2012.

- [42] P. Cappers, “Mass market demand response and variable generation integration issues: A scoping study,” 2012.
- [43] M. Hummon, D. Palchak, P. Denholm, J. Jorgenson, D. J. Olsen, S. Kiliccote, N. Matson, M. Sohn, C. Rose, J. Dudley *et al.*, “Grid integration of aggregated demand response, part 2: Modeling demand response in a production cost model,” 2013.
- [44] H. Aalami, M. P. Moghaddam, and G. Yousefi, “Demand response modeling considering interruptible/curtailable loads and capacity market programs,” *Applied Energy*, vol. 87, no. 1, pp. 243–250, 2010.
- [45] M. E. Dyson, S. D. Borgeson, M. D. Tabone, and D. S. Callaway, “Using smart meter data to estimate demand response potential, with application to solar energy integration,” *Energy Policy*, no. 0, pp. –, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0301421514003681>
- [46] A. Rabl, “Parameter estimation in buildings: methods for dynamic analysis of measured energy use,” *Journal of Solar Energy Engineering*, vol. 110, no. 1, pp. 52–66, 1988.
- [47] T. Facchinetti and M. Della Vedova, “Real-time modeling for direct load control in cyber-physical power systems,” *Industrial Informatics, IEEE Transactions on*, vol. 7, no. 4, pp. 689–698, Nov. 2011.
- [48] J. Mathieu and D. Callaway, “State estimation and control of heterogeneous thermostatically controlled loads for load following,” in *System Science (HICSS), 2012 45th Hawaii International Conference on*, jan. 2012, pp. 2002–2011.

- [49] M. Kamgarpour, C. Ellen, S. Gerwinn, N. Mathieu J. L. Mullner, and J. Lygeros, “Modeling options for demand side participation of thermostatically controlled loads,” in *Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP)*, IEEE, 2013, pp. 1–15.
- [50] ERCOT, “Requirements for aggregate load resource participation in the ERCOT markets,” 2014. [Online]. Available: <http://www.ercot.com/content/services/programs/load/laar/RequirementsforAggregateLoadResourceParticipationinth.zip>
- [51] J. MacDonald, P. Cappers, D. Callaway, and S. Kiliccote, “Demand response providing ancillary services,” *Presented at Grid-Interop*, 2012.
- [52] M. Wytock and J. Z. Kolter, “Contextually supervised source separation with application to energy disaggregation,” *arXiv preprint arXiv:1312.5023*, 2013.
- [53] D. G. Robertson, J. H. Lee, and J. B. Rawlings, “A moving horizon-based approach for least-squares estimation,” *AIChE Journal*, vol. 42, no. 8, pp. 2209–2224, 1996.
- [54] E. L. Haseltine and J. B. Rawlings, “Critical evaluation of extended kalman filtering and moving-horizon estimation,” *Industrial & Engineering Chemistry Research*, vol. 44, no. 8, pp. 2451–2460, 2005. [Online]. Available: <http://pubs.acs.org/doi/abs/10.1021/ie034308l>
- [55] C. V. Rao, J. B. Rawlings, and J. H. Lee, “Constrained linear state estimation—A moving horizon approach,” *Automatica*, vol. 37, no. 10, pp.

- 1619 – 1628, 2001. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0005109801001157>
- [56] G. Welch and G. Bishop, “An introduction to the kalman filter,” UNC-CH Computer Science, Tech. Rep. 95041, 1995.
- [57] S. Kullback, *Information Theory and Statistics*. Mineola, New York: Dover Publications Inc, 1968.
- [58] N. A. Smith and R. W. Tromble, “Sampling uniformly from the unit simplex,” *Johns Hopkins University, Tech. Rep*, 2004.
- [59] C. M. Bishop *et al.*, *Pattern recognition and machine learning*. springer New York, 2006, vol. 1.
- [60] M. Duvall, E. Knipping, M. Alexander, L. Tonachel, and C. Clark, “Environmental assessment of plug-in hybrid electric vehicles,” *EPRI, July*, 2007.
- [61] J. Lopes, F. Soares, and P. Almeida, “Integration of electric vehicles in the electric power system,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 168–183, Jan 2011.
- [62] M. D. Galus, M. G. Vaynskiy, T. Krause, and G. Andersson, “The role of electric vehicles in smart grids,” *Wiley Interdisciplinary Reviews: Energy and Environment*, vol. 2, no. 4, pp. 384–400, 2013. [Online]. Available: <http://dx.doi.org/10.1002/wene.56>
- [63] C. Guille and G. Gross, “A conceptual framework for the vehicle-to-grid implementation,” *Energy Policy*, vol. 37, no. 11, pp. 4379–4390, 2009.

- [64] A. Brooks, T. Gage, and A. Propulsion, "Integration of electric drive vehicles with the electric power grid—a new value stream," in *18th International Electric Vehicle Symposium and Exhibition, Berlin, Germany*. Citeseer, 2001, pp. 20–24.
- [65] P. Sanchez-Martin, G. Sanchez, and G. Morales-Espana, "Direct load control decision model for aggregated ev charging points," *Power Systems, IEEE Transactions on*, vol. 27, no. 3, pp. 1577–1584, Aug 2012.
- [66] W. Su and M.-Y. Chow, "Performance evaluation of an eda-based large-scale plug-in hybrid electric vehicle charging algorithm," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 308–315, 2012.
- [67] Y. He, B. Venkatesh, and L. Guan, "Optimal scheduling for charging and discharging of electric vehicles," *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1095–1105, 2012.
- [68] N. Rotering and M. Ilic, "Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets," *Power Systems, IEEE Transactions on*, vol. 26, no. 3, pp. 1021–1029, Aug 2011.
- [69] J. Tomić and W. Kempton, "Using fleets of electric-drive vehicles for grid support," *Journal of Power Sources*, vol. 168, no. 2, pp. 459–468, 2007.
- [70] W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre, "A test of vehicle-to-grid (v2g) for energy storage and frequency regulation in the pjm system," *Results from an Industry-University Research Partnership*, p. 32, 2008.

- [71] P. Finn, C. Fitzpatrick, and D. Connolly, "Demand side management of electric car charging: Benefits for consumer and grid," *Energy*, vol. 42, no. 1, pp. 358 – 363, 2012, 8th World Energy System Conference, {WESC} 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0360544212002435>
- [72] D. Dallinger, D. Krampe, and M. Wietschel, "Vehicle-to-grid regulation reserves based on a dynamic simulation of mobility behavior," *Smart Grid, IEEE Transactions on*, vol. 2, no. 2, pp. 302–313, June 2011.
- [73] O. Sundstrom and C. Binding, "Flexible charging optimization for electric vehicles considering distribution grid constraints," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 26–37, 2012.
- [74] B. Dietz, K. Ahlert, A. Schuller, and C. Weinhardt, "Economic benchmark of charging strategies for battery electric vehicles," in *PowerTech, 2011 IEEE Trondheim*. IEEE, 2011, pp. 1–8.
- [75] S. Shao, M. Pipattanasomporn, and S. Rahman, "Grid integration of electric vehicles and demand response with customer choice," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 543–550, March 2012.
- [76] P. S. Hu and T. R. Reuscher, "Summary of travel trends: 2001 national household travel survey," 2004.
- [77] C. Binding, D. Gantenbein, B. Jansen, O. Sundstrom, P. Andersen, F. Marra, B. Poulsen, and C. Traeholt, "Electric vehicle fleet integration in the danish edison project - a virtual power plant on the island of bornholm," in *Power and Energy Society General Meeting, 2010 IEEE*, July 2010, pp. 1–8.

- [78] E. Knipping and M. Duvall, “Environmental assessment of plug-in hybrid electric vehicles, volume 2: United states air quality analysis based on aeo-2006 assumptions for 2030,” *Electric Power Research Institute and Natural Resources Defense Council*, 2007.
- [79] S. Deilami, A. S. Masoum, P. S. Moses, and M. A. Masoum, “Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile,” *Smart Grid, IEEE Transactions on*, vol. 2, no. 3, pp. 456–467, 2011.
- [80] M. Shaaban, M. Ismail, E. El-Saadany, and W. Zhuang, “Real-time pev charging/discharging coordination in smart distribution systems,” *Smart Grid, IEEE Transactions on*, vol. 5, no. 4, pp. 1797–1807, July 2014.
- [81] F. Marra, G. Y. Yang, C. Traholt, E. Larsen, C. N. Rasmussen, and S. You, “Demand profile study of battery electric vehicle under different charging options,” in *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–7.
- [82] G. Optimization, “Inc.: Gurobi optimizer reference manual (2012),” 2014.