Modeling and Control of Electric Loads for Ancillary Services and Decarbonization

by

Oluwagbemileke Oyefeso

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical and Computer Engineering)
in The University of Michigan
2025

Doctoral Committee:

Associate Professor Johanna L. Mathieu, Co-Chair
Professor Ian A. Hiskens, Co-Chair
Associate Professor Mads Almassakhi
Associate Professor Christian Lastoskie

Oluwagbemileke Oyefeso
ooyefeso@umich.edu
ORCID iD: 0000-0002-6400-3679

© Oluwagbemileke Oyefeso 2025

Dedication

To my amazing parents, Gbenga and Nike.

To my brothers and sister, the best siblings one could hope for, Oluwapelumi, Oluwatunmise and Oluwamayokun.

To my loving partner, Nicole Wiggins.

To all who have been a part of my journey.

To those dreaming of the seemingly impossible, keep at it.

ACKNOWLEDGEMENTS

Firstly, thanks to my research advisors, Professors Johanna Mathieu and Ian Hiskens.

I am grateful to Johanna for being an incredible mentor and a steadfast support system throughout my PhD journey. Your unwavering belief in me, often greater than my own, carried me through challenging times and inspired me to persevere. Your persistent push for excellence, coupled with an unmatched attention to detail, has not only made me a better researcher but also a better person. Johanna's genuine care for her students as individuals, not just as researchers, is a quality I deeply admire and aspire to emulate. You have been a mentor in the truest sense, someone I look up to, someone I hope to emulate, and someone whose example I will carry forward in my own journey. May I meet more extraordinary individuals like you, and may I someday raise and inspire others as you have inspired me. Thank you, Johanna, for everything.

To Ian, I would like to express my heartfelt gratitude for being an unwavering source of support and guidance throughout my academic journey. Your constructive feedback has profoundly shaped me into a better researcher. Whenever we had to navigate the challenges of the 15-hour time difference between Australia and the United States, you were flexible and always made time for me, demonstrating your dedication. Beyond academics, your kindness and generosity, opening your home to us for get-togethers and Thanksgiving, making us feel like family, provided a sense of home away from home and was cherished. Thank you for everything, Ian.

To Professor Mads Almassalkhi, I extend my sincere gratitude for your invaluable support and encouragement. As a co-author and collaborator, your insights and

expertise enriched my work. Your constructive feedback and thoughtful guidance have been instrumental in shaping my research, and your steady encouragement has been a source of motivation during pivotal moments. I deeply appreciate your willingness to share your time and knowledge, and your support as a member of my dissertation committee has meant a great deal to me. Thank you, Mads, for your contributions and for being such an inspiring collaborator.

To Professor Christian Lastoskie, thank you for being a part of my dissertation committee. Your support, inputs and guidance are truly appreciated.

To Dr. Ioannis Granitsas, thank you for being an exceptional colleague, collaborator, and friend throughout my PhD journey. Sharing the lab and working on the same project for over four years has been an incredible experience, and I have learned a lot from your expertise, dedication, and collaborative spirit. Co-authoring two papers, three technical reports and one open-source code repository together is a testament to the synergy we built, and your insights and work ethic have consistently elevated our research. Beyond academics, Ioannis has been a constant source of support, camaraderie, and inspiration, making this long journey an enjoyable one. Thank you, Ioannis, for everything, it has been an honor to work alongside you.

To Dr. Drew Geller, I express my heartfelt thanks for your unwavering support and collaboration. As a co-author and collaborator, your insightful contributions and thoughtful feedback have been invaluable. Thank you, Drew, for leading the efforts on our R & D 100 award, and for being an integral part of my journey.

I am deeply thankful to Dr. Bosun Roy-Layinde, my roommate, friend, and steadfast pillar of support over the past five years. Your unwavering encouragement and camaraderie have played an invaluable role on this scholarly odyssey. From teaching me to drive to being a shoulder to lean on during challenging times, you have consistently been there, making the highs more joyful and the lows more bearable. Bosun, your friendship has been a source of strength and inspiration, and I am profoundly grateful to have had you by my side throughout this journey. Thank you for everything and I look forward to our future celebrations and endeavors!

To my best friend of almost a decade, Segun Ibiyemi, sharing the highs and lows of our PhD experiences has been both grounding and uplifting, and your friendship

forever remains a constant source of strength. Your home in Atlanta has always been a haven for me, offering warmth and belonging whenever I visit. Thank you for being an incredible friend, a trusted confidant, and a steadfast cheerleader. I am truly fortunate to have you in my corner, and I look forward to many more years of friendship and shared successes.

To my amazing parents, I owe everything to you. Without your tireless sacrifices, unwavering prayers, and unconditional love, I would not be the person I am today. You have shaped me with your guidance, wisdom, and values, and it is a privilege to call you my parents. Your steadfast commitment to my progress, always rooting for my success and happiness, is a testament to the depth of your love. Your strength and sacrifices are the foundation of all I've achieved. Thank you for being my greatest champions, my unwavering support system, and simply the best parents anyone could ever hope for. I am endlessly grateful for you. For my mother, just like Damini Ogulu said, "you can never take the pride away from my mama. You'd be proud of me too, if you knew where I was coming from."

To my incredible siblings, the best anyone could ever hope for, thank you for being my north star, my grounding force, and my constant source of inspiration. You make life worth living and are the reason I keep pushing forward, even in the toughest moments. Your unwavering love, support, and belief in me have been the foundation of everything I have achieved. Without you, there is no me, and I am endlessly grateful for the bond we share. You are, and always will be, my greatest blessing.

To my loving partner, Nicole Wiggins, words cannot fully express how grateful I am for your unwavering support, love, and encouragement over the past few years. You have been my pillar, offering strength, understanding, and comfort through every challenge and triumph. Your presence in my life has made this journey more meaningful, and you have truly made the PhD only the second most important thing I leave Michigan with. Thank you for being my rock, my partner, and cheerleader. I am so fortunate to have you by my side and I am excited to see what the future holds for us.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF FIGURES x

LIST OF TABLES xii

LIST OF ABBREVIATIONS xiii

ABSTRACT xv

CHAPTER

1. Introduction 1

1.1 From Demand Response to Virtual Power Plants 2
1.2 Electric Load Control 4
1.3 Open Problems 10
1.4 Organization and Contributions 12

2. Control of Aggregate Air-Conditioning Load using Packetized Energy Concepts 15

2.1 Notation 16
2.2 Chapter Introduction 17
2.3 Air Conditioner Control with PEM 19
2.3.1 AC Overview and Model 19
2.3.2 Packetized Energy Management Overview 21

2.3.3 Air Conditioners as a Packetized Load 24
2.4 Controller Extensions 25
2.5 Case Studies 28
2.5.1 Simulation Scenarios and Parameters 28
2.5.2 Simulation Results 30
2.6 Conclusion 34

3. Device-Driven Modeling and Control of Thermostatic Loads: A Bottom-Up Approach 36

3.1 Chapter Introduction 37
3.2 Framework 39
3.3 Device-Driven Load Control for Air Conditioners 41
3.3.1 Individual AC Model 41
3.3.2 Extended PEM Overview and Shortcomings 43
3.3.3 New Device-Driven Load Control Strategy 46
3.4 Aggregate Modeling 51
3.5 Case Study 58
3.5.1 Simulation Scenarios and Parameters 58
3.5.2 Simulation Results 62
3.5.3 Aggregate Model Results 65
3.6 Summary and Conclusion 67

4. A Hardware-in-the-Loop Experimental Testbed using Air Conditioners for Grid Balancing 68

4.1 Chapter Introduction 69
4.2 Hardware-in-the-loop Architecture 71
4.2.1 Experimental Testbed Design 71
4.2.2 Simulation Testbed Design 75
4.2.3 Aggregator's Load Control Algorithms 76
4.2.4 Data Acquisition and Control System 76
4.2.5 Communication System 77
4.3 HIL Testbed Validation 78
4.3.1 Experimental Testbed Validation 78
4.3.2 Simulation Testbed Validation 80
4.4 Hardware-in-the-Loop Experiments 82
4.4.1 Setup 82

4.4.2 Results and Discussion 85
4.5 Conclusions 90

5. Comparative Evaluation of Energy Modeling Methods in Residential Electrification and Decarbonization 91

5.1 Introduction 92
5.1.1 Aims and Background 92
5.1.2 Motivation and Contributions 93
5.1.3 Chapter Organization 95
5.2 Literature Review 95
5.2.1 System-Level Energy Models 95
5.2.2 Tailored Energy Models 97
5.3 Energy Modeling Frameworks – Residential Electrification . . 98
5.3.1 Simple Calculators 98
5.3.2 Thermodynamic Models 101
5.3.3 Load Shape Modeling 103
5.3.4 Building Simulation 104
5.4 Proposed Methodology for Comparison across Energy Modeling Frameworks 107
5.4.1 Data to Be Gathered from PSI Homes 107
5.4.2 Framework 1: Simple Calculators 108
5.4.3 Framework 2: Thermodynamic Model 108
5.4.4 Framework 3: Load Shape Modeling 109
5.4.5 Framework 4: Building Simulation 110
5.4.6 Comparison Metrics and Forecasting Procedure . . . 110
5.4.7 Outlook for Empirical Validation 111
5.5 Discussion 113
5.5.1 Comparison of Modeling Frameworks 113
5.5.2 Implications for Stakeholders 114
5.5.3 Balancing Accuracy and Practicality 115
5.5.4 Trends and Challenges 115
5.6 Conclusion 116

6. Conclusion 117

6.1 Dissertation Summary and Contributions 117
6.2 Possible Future Research Extensions 118

APPENDIX 119

A. Communication Architecture enabling Cyber-Physical Load Control for Frequency Regulation 120

BIBLIOGRAPHY 123

LIST OF FIGURES

Figure

2.1	PEM Scheme for TCLs.	23
2.2	Extended PEM for closed-loop control of an AC aggregation.	27
2.3	Reference tracking, resource availability, and flexibility for all cases.	31
2.4	Effect of lock-out and controller extensions on AC availability.	32
2.5	Close-up of the tracking results for Cases 1 and 2.	33
2.6	Extended PEM providing 1 MW frequency regulation, with new M_{on}. PJM composite score is 0.84.	35
3.1	Device-driven load control architecture and flexibility across the different strategies.	40
3.2	Left Plot: Correlation between temperature and elapsed epoch. Right Plot: Log Plot showing correlation between temperature-based turn-off rate function $\mu^{\text{off}}(T_{\text{a},n}[k])$ and time-based turn-off rate function $\gamma(t_n[k])$.	49
3.3	Turn-Off Request Probability computation using 2 rate functions ($\mu^{\text{off}}(T_{\text{a},n}[k])$ and $\gamma(t_n[k])$) vs 1 rate function ($\mu^{\text{off}}(T_{\text{a},n}[k])$).	50
3.4	Temperature-based and time-based turn-off probability evolution with temperature.	50
3.5	State bin transition model for ACs under the new device-driven controller.	53
3.6	State bin transition model for ACs under the extended PEM scheme.	57
3.7	Top 3 plots: reference tracking for cases 4B, 5B and 6B respectively; bottom plot: case 6B with perfect communication network.	63
3.8	Case 1B: simulation vs 3-steps ahead aggregate model prediction.	66
4.1	Hardware-in-the-Loop experiment architecture.	72

4.2 Left: Model houses are constructed inside a four foot cubic foam box. The internal heat source consists of a hydronic loop with an electric water heater. A duct fan forces air through the heat exchanger and mixes the room air in addition to the AC's fan. Right: The 20 units are stacked in pallet racks, minimizing cable lengths and overall footprint. 73
4.3 Fractional power variation when the ACs are on for the experimental ACs (right) and real ACs (left). The distributions approximately match. 79
4.4 Active (left) and reactive (right) power to temperature correlation plots for the real ACs (top) and the experimental ACs (bottom). Note the difference in x-axis scale of the histograms for reactive power. 79
4.5 Active (left) and reactive (right) power comparison between virtual and experimental ACs. 81
4.6 Comparison of controller tracking performance in Case 1. 86
4.7 PEM controller tracking performance in Case 3 (imperfect communication network and larger amplitude reference signal than Case 1). 87
4.8 Transformer power flows in case 1 with the PEM controller. The transformer ratings were 1 p.u. 88
4.9 Percentage of experimental and virtual ACs that are locked (top), on (middle), and off (bottom) for the PEM controller in Case 1. . . 89
A.1 Architecture of the Cyber-Physical System 121
A.2 Communication Architecture between Field Testbed and Controller. 121

LIST OF TABLES

Table

1.1	Literature Survey of TCL Control Strategies	7
2.1	Simulation Parameters	28
2.2	Controller performance results across all cases	29
2.3	Average AC availability	32
3.1	System Architectures across device-driven load controllers.	40
3.2	Simulation Conditions	59
3.3	Simulation Parameters	59
3.4	Controller performance results across nominal cases	60
3.5	Controller performance results across extreme cases	61
3.6	Model Prediction Accuracy	66
4.1	Simulation Conditions	82
4.2	HIL Experiment Results for the 3 controllers across the 10 experiments. Nom: nominal, Ext: extreme condition.	84
5.1	Comparison of Energy Modeling Frameworks	112

LIST OF ABBREVIATIONS

AC Air Conditioner

AGC Automatic Generation Control

DAQ Data Acquisition

DER Distributed Energy Resource

DERMS Distributed Energy Resource Management Systems

DOE Department of Energy

DSM Demand Side Management

ESS Energy Storage System

ETP Equivalent Thermal Parameter

EV Electric Vehicle

FERC Federal Energy Regulatory Commission

HDD Heating Degree Days

HIL Hardware-in-the-Loop

HVAC Heating, Ventilation and Air Conditioning

Hz Hertz

IPCC Intergovernmental Panel on Climate Change

kW Kilowatt

kWh Kilowatt-hour

MAPE Mean Absolute Percentage Error

MQTT Message Queuing Telemetry Transport

MW Megawatt

NRMSE Normalized Root Mean Square Error

PEM Packetized Energy Management

PI Proportional Integral

PJM Pennsylvania-New Jersey-Maryland Interconnection

RegD Regulation-D

RMS Root Mean Square

TCL Thermostatically Controlled Load

TCP Transmission Control Protocol

US United States

USD United States Dollar

VPP Virtual Power Plant

ABSTRACT

Towards decarbonizing power systems and achieving sustainability goals, traditional power plants, such as coal fired power plants, are being phased out and replaced with renewable generation alternatives. Historically, synchronous generators of traditional power plants provided much of the essential load balancing services required to match scheduled generation and demand, ensuring power system reliability. Hence, with the sunsetting of traditional generation, new sources of load balancing flexibility are needed. Furthermore, the inherent intermittent and variable nature of renewable generation typically requires an increase in the load balancing services needed. Aggregations of electric loads, such as thermostatically controlled loads (TCLs), hold substantial potential to offer this flexibility, given their thermal inertia and ubiquity.

TCLs are an ubiquitous resource with inherent temporal flexibility that allows them cycle on/off and consume power in cycles. This flexibility can be leveraged for power system objectives. The abundance and spatial distribution of TCLs make them a prime and valuable candidate to provide flexibility to the power grid.

This dissertation develops models and control algorithms for leveraging the flexibility of TCL resources and practically demonstrates the ancillary service and decarbonization potential of these loads. Building on established hierarchical load control methods, a novel device-driven approach for coordinating aggregations of TCLs is created. The developed control design was tested and validated via simulations and hardware-in-the-loop (HIL) experiments. Also, communication architectures enabling cyber-physical demonstrations of load control for frequency regulation were developed.

In this dissertation, the state-of-the-art device-driven packetized energy management (PEM) aggregate load control strategy, which coordinates only thermostatic

loads without compressors, is improved upon to create a new device-driven method that enables coordination and control of all classes of thermostatic loads. The new control approach when compared to the state-of-the-art PEM, significantly improves control flexibility and signal tracking performance with the load aggregation. It is demonstrated with simulation and HIL experiments that the newly developed device-driven control is able to satisfactorily (according to industry standards) provide up to 1 MW capacity of frequency regulation balancing service using about a thousand residential air conditioners. An aggregate model of the new device-driven strategy that can be used for predictive control and analyzing system dynamics is also developed.

There exists significant efforts to decarbonize by replacing fossil-fuel dependent systems like space heating with electric heat pumps, appliances such as dryers and water heaters, with electric versions. The transition to electrified homes is crucial for climate goals but requires precise evaluation methods to understand energy, economic, and environmental impacts. A systematic review of energy modeling approaches, used in electrification and decarbonization impact assessments, is also presented in this dissertation.

Overall, this dissertation attempts to establish credibility for aggregate control of electric loads in providing power system balancing services and harnessing the decarbonization potential of load electrification. This dissertation identifies and overcomes challenges involving electric load aggregations, showing they can indeed provide the additional flexibility required in modern power systems.

CHAPTER 1

Introduction

A fundamental objective in power system operations is matching the scheduled energy generation with the energy demand. A large enough imbalance between demand and scheduled supply could lead to issues with the reliability and stability of the power system [1]. Therefore, there exist many strategies in place to minimize the imbalance between the power and energy demand and scheduled supply. Strategies such as Automatic Generation Control (AGC) and droop speed control take action on the supply side to vary the scheduled power generation to match demand [2]. There are also approaches for modifying demand, this is referred to as Demand Side Management. Demand Side Management (DSM) plays a crucial role in ensuring the reliability and efficiency of power systems. It involves a set of actions undertaken on the consumer side to modify the electricity consumption pattern, responding to the needs of the power grid [3]. DSM also includes energy-efficiency and conservation voltage reduction strategies. Historically, the supply side strategies played a greater role in minimizing the imbalance between demand and supply. However, increased renewable energy penetration on the supply side introduces higher levels of uncertainty and variability in power generation (due to the intermittent nature of renewable energy), accentuating the importance of DSM in balancing supply and demand [2]. This means additional DSM alternatives to minimizing the supply/demand imbalance are needed. This has given rise to the advent of Virtual Power Plants (VPPs) which refer to an aggregation of Distributed Energy

Resources (DERs) being controlled for power system operational objectives. VPPs not only seek to ensure power system reliability by providing an alternative means to maintain power balance, but could also be pivotal in decarbonizing the power sector by providing renewables with the balancing support they need. The work presented in this dissertation focuses on aggregations of distributed energy resources, particularly flexible loads, to offer balancing services to the power grid. We focus on developing electric load control and coordination strategies, practical demonstrations of these control designs and detailed performance and controller assessments. In the following sections, a brief background and literature review on electric load control is provided, then the organization structure and contributions of the entire dissertation is outlined.

1.1 From Demand Response to Virtual Power Plants

Demand Side Management (DSM) encompasses both direct load control and behavioural modification techniques targeted at modifying consumer demand such as financial incentives, consumer education and other such strategies. DSM includes demand response, energy-efficiency and conservation voltage reduction strategies. Demand response traditionally refers to load curtailment during hours of peak electricity consumption [3]. The integration of DSM into power systems has evolved over the years, from traditional demand response mechanisms to the contemporary concept of VPPs, a dispatchable aggregation of DERs. Looking at the historical development of demand response, the rise of VPPs, their relevance in the era of increased renewable energy penetration, and their potential in contributing to the decarbonization of power systems, we see that demand response has witnessed a transformation from manual load shedding to sophisticated, technology-driven approaches.

Early demand response initiatives date back to the nascent U.S. electric power industry of the 1880s, where electricity rate structures were used to encourage shifting electricity usage to off-peak hours [4]. This laid the foundation for future time-of-use and price-based demand response mechanisms. However, various events such as

the California electricity crisis, still make regulators skeptical about real-time pricing demand response as it often exposes customers to significant price volatility [5]. The rapid demand growth of the mid-20th century, driven by industrialization and urbanization, necessitated the use of load shedding. Load shedding is one of the earliest and most basic forms of demand response; it involves direct load curtailment by disconnecting portions of the power network, in order to prevent blackouts when generation simply can't match demand. Load shedding was employed during the energy crisis of the 1970s and California electricity crisis of the early 2000s [6]. It remains available till date as a last resort technique to preventing blackouts due to extreme weather events or other emergency scenarios. References [7–9] discus work done in the mid 20th century to automate load shedding, paving the way for interruptible load programs where customers were rewarded for curtailing their electricity consumption during peak demand hours. The several energy crisis and unpredictable demand patterns birthed many energy efficiency programs where utilities provided incentives to consumers to adopt energy-efficient appliances [10].

Electricity market restructuring from vertically integrated utilities to the system based on economic markets made integration of demand response into power markets desirable. Eliminating barriers to integrating demand response into the electric power markets was a primary goal of the 2005 Energy Policy Act and the 2007 Energy Independence Security Act. The 2009 American Recovery and Reinvestment Act kick started financing for the modernization of the electric grid - the "smart grid" movement [11]. These policies advanced telemetry, communication and control technologies in the power sector via smart meters and other modernization initiatives, leading to the development of automated systems that could respond to signals from the grid, paving the way for more dynamic and responsive demand-side interventions [12]. This signaled a new paradigm in demand response popularly termed demand dispatch. "Demand dispatch involves the capability to aggregate and precisely control loads on command; unlike traditional demand response, demand dispatch is active and can be deployed at all times and not just at peak times" [3]. Essentially, it involves transforming loads from a passive demand resource which simply 'responds' to commands to an active 'dispatchable' power system resource.

With the proliferation of DERs and the increased penetration of renewable energy sources, the concept of VPPs emerged. VPPs refer to a portfolio of actively controlled DERs for power system objectives [13]. VPPs leverage smart technologies to aggregate and manage a diverse array of distributed energy assets, creating a flexible and responsive virtual power entity. Recent economic policies motivating this are FERC Orders 719, 745, 755, 784 and 2222. Order 719 made several reforms to further eliminate barriers to demand response participation in power markets [14], while Order 745 addressed compensation for demand response participation [15]. Order 755 ensures fair access to the ancillary service markets for these resources along with equitable compensation [16]. Order 784 stipulates better compensation for faster response to ancillary service requests. Order 2222 requires grid operators revise their rules to enable aggregations of DERs to participate alongside other traditional resources in the wholesale markets [17, 18].

Beyond power system reliability, VPPs offer cost savings and environmental benefits. The US Department of Energy estimates a current national VPP capacity of up to 60GW [19], with potential resource adequacy savings estimates of several billion dollars [13]. Utilizing these VPPs for energy balancing and ancillary service offerings (e.g. frequency regulation and reserves) would offset several tonnes of CO_2 [19] providing billions of dollars in societal benefit [20]. This is important as the Electric Power sector accounts for 25% of U.S. Green House Gas emissions [21]. The decarbonization potential of VPPs is a critical aspect in achieving sustainability goals and addressing climate change challenges.

This transition from individual demand response actions to the orchestration of a virtual power system is a testament to the constantly evolving landscape of demand side management. In this dissertation, we are particularly interested in the electric loads subcategory of VPP resources.

1.2 Electric Load Control

Electric Load Control involves managing electricity consumption at the consumer level by actively controlling flexible consumer loads such as heating, ventilation and

air conditioning (HVAC) equipment, refrigerators, plug-in electric vehicles and battery energy storage systems. Research shows these flexible loads can be actively controlled to provide power system services at similar scale and with similar levels of performance guarantees as conventional generation methods [2, 13, 22, 23]. References [24–27] discuss methods of flexibly controlling commercial and industrial HVAC, [28–31] discuss controlling residential thermostatically controlled loads (TCLs) such as air conditioners, heat pumps, refrigerators and water heaters for power system objectives, [32, 33] considers plug-in electric vehicles while [34–36] discuss energy storage systems providing ancillary services for power systems.

Residential loads make up 39% of the overall electricity consumption in the United States, commercial loads make up 35% and industrial loads make up 26% [37]. Over 50% of residential energy consumption is by flexible loads with 19% being air conditioning, 12% space heating, 12% water heating and 8% refrigerators [38]. In the commercial sector, about 30% of the electricity consumption is by flexible heating and cooling loads [39]. With transportation, electric vehicle adoption rate soars continually, with over 2 million electric vehicles currently registered in the U.S. [40] and an average average light-duty EV battery ranges from 15-100 kWh [41]. These statistics show that the importance of flexible load control can not be overstated.

The dynamics of commercial loads and residential loads differ and so do the control strategies. In this dissertation, we shall be focusing on residential thermostatic load control. As stated earlier, more than half of the residential energy consumption can be flexibly controlled to aid power system objectives. Furthermore, over 90% of US homes contain TCLs such as air-conditioners, water heaters and refrigerators, making this a valuable ubiquitous resource [42]. The temporal flexibility of these TCLs as a result of their inherent thermal inertia and their distributed geographic locations make them prime and valuable candidates for providing flexibility to the power grid. In summary, Demand Side Management, particularly through the lens of Electric Load Control VPP offering, could emerge as a cornerstone for ensuring the reliability, efficiency, and sustainability of modern power systems.

There exist numerous load modeling [43–46], estimation [47–50], parameter & system identification [51, 52] and control techniques [6, 53, 54] that underpin the ther-

mostatic load control literature. However, in this work we shall focus particularly on the load modeling & controller design, and practical demonstrations of this technology. The thermostatic load control approaches in the literature can be grouped and categorized according to their system architectures, communication requirements, consumer impact, as well as actuation methods, which is shown in Table 1.1. Each approach has its merits and challenges, understanding these help clarify where existing methods fall short and how emerging approaches, like those explored in this dissertation, can offer a more effective solution. Table 1.1 also provides my subjective assessment on the practical scalability of the various strategies, informed by insights from pilot demonstrations in the literature [55, 56], modeling complexity required, and data & communication requirements for each approach.

From Table 1.1, we see based on control strategy, [31, 46, 57–60] adopt a Markov-model based probabilistic switching approach wherein a set of switching probabilities are periodically broadcasted to TCLs which then make ON/OFF switching decisions according to these probabilities. This strategy typically employs a centralized control architecture with a top-down, unidirectional communication approach, i.e., the coordinating entity e.g. third-party aggregator or utility is "at the top" of the control hierarchy and sends "down" control actions to the TCLs. Typically, this control approach is non-disruptive to consumer thermal comfort, that is, the quality of service provided by the loads to the consumer remains unchanged; the loads provide their primary heating or cooling function unaffected by the secondary function - an actively controlled DER for power system objectives.

References [30, 61–63] in Table 1.1 utilize a distributed, device-driven, bottom-up energy packetization approach. Here, devices dictate when they consume energy based on their local state information. Communication is bi-directional because devices send energy consumption requests to the coordinating entity who then responds with an approval/denial. This is referred to as "bottom-up" since control actions are primarily driven upwards from the horde of devices to a coordinating aggregator. It is considered a 'distributed' approach since it entails both decentralized control decisions (individual local device requests) and a centralized aggregator approval/denial action. It is also typically non-disruptive of consumer thermal comfort.

Table 1.1: Literature Survey of TCL Control Strategies

Control Strategy	Architecture	Hierarchy	Communication Mode	Consumer Impact	Actuation Mode	Scalability	Algorithms	Reference
Probabilistic Switching	Aggregator-driven	Top-down, Centralized	Uni-directional	Non-disruptive	ON/OFF switching	Feasible	Markov-Model based	[6, 31, 46, 57–60]
Packetized Energy	Device-driven	Bottom-up, Distributed	Bi-directional	Non-disruptive	ON/OFF switching	Feasible	Energy packetization	[30, 61–63]
Set-point Variation	Aggregator-driven	Top-down, Centralized	Uni-directional	Disruptive	Set-point varying	Feasible	Set-point shifting	[64–68]
Online Convex Optimization	Aggregator-driven	Top-down, Centralized	Bi-directional	Disruptive	ON/OFF switching	Challenging	convex online gradient descent	[69–74]
Reinforced Learning	Aggregator-driven	Top-down, Centralized,	Bi-directional	Disruptive	ON/OFF switching	Challenging	Reinforced Learning	[75–77]
Distributed-consensus	Device-driven	Bottom-up, Distributed	Bi-directional	Disruptive	Varies	Challenging	Ratio consensus, Gossip, Constrained Fair-Splitting	[78–81]
Extremum-seeking Control	Aggregator-driven	Top-down, Centralized	Uni-directional	Disruptive	Set-point varying	Challenging	Pertubation-based	[82–85]

References [64–68] of Table 1.1 have similar system architectures to the probabilistic switching approach discussed above but involve adjusting the operational set-point of TCLs (e.g., raising the temperature of TCLs) to increase or decrease energy consumption. While effective in managing load during peak periods, this approach is disruptive to users since it directly affects their comfort levels. This strategy, while useful in specific scenarios, does not align well with the goals of this dissertation, which seeks non-disruptive, scalable solutions that maintain consumer comfort. The research here aims to leverage strategies that allow for seamless integration into existing systems without compromising user-preferred settings.

Table 1.1 also highlights recent sophisticated approaches that use online convex optimization and reinforcement learning, to improve TCL control. In [69–74], an online convex optimization approach based on binary convex online gradient descent is used for TCL control. It is an aggregator-driven approach with the aggregator solving optimization problems like minimizing power deviation or energy costs in real-time, considering current state information and system operator command. Likewise, the reinforced learning approach of [75–77] where the aggregator uses reinforced learning algorithms to learn optimal control policies based on the feedback received from the devices and the grid. However, the complexity of these algorithms and their need for significant computational resources can hinder practical scalability. Moreover, their centralized, top-down nature means they can still face issues related to communication bottlenecks and reduced flexibility. These advanced techniques underscore the shift towards smarter, adaptive control, but their frameworks do not fully solve the issues of scalability and adaptability, [55] shows it is often disruptive and difficult to scale in practice, when compared with alternative strategies. This dissertation builds on the principles of adaptability seen here but seeks to incorporate them into a distributed, device-driven model that avoids the above mentioned pitfalls.

References [78–81] of Table 1.1 take a distributed-consensus strategy based on ratio consensus, gossip and constrained fair-splitting, for DER demand dispatch. It employs a bi-directional, distributed communication & control approach. Scaling these methods remains a challenge due to the computational and communication overhead involved. However, the emphasis on device-driven control aligns closely

with this dissertation's aim to explore bottom-up approaches that improve scalability and practical feasibility without compromising on responsiveness or flexibility.

The extremum-seeking control approaches [82–85] shown in Table 1.1 offer an adaptive optimization technique that adjusts system settings by introducing small perturbations and analyzing the response to find optimal operating points. By applying small perturbations and analyzing the system's response, the controller iteratively adjusts settings to find the optimal operation, akin to gradient descent but without direct knowledge of the gradient. With TCL control, the extremum-seeking approach attempts to dynamically optimize energy consumption or cost without requiring a detailed system model, making it adaptable to nonlinear, variable conditions. On the other hand, the need for continuous perturbation can disrupt normal operations, affecting user comfort. Additionally, the centralized nature of traditional extremum-seeking implementations introduces communication bottlenecks and limits scalability, especially when managing large, distributed systems. While extremum-seeking offers robust adaptive control, its challenges in scalability and consumer impact highlight the need for more device-driven distributed solutions. This dissertation seeks to retain adaptability within a distributed, device-driven framework.

The control strategies described above and in Table 1.1 each offer varying degrees of scalability, flexibility, and consumer impact. However, many existing methods rely on aggregator-driven, centralized, top-down architectures, which despite their centralized visibility and controllability advantages, can limit their practicality. The practical challenges arise from the fact that these methods require the coordinating entity to collate state information of loads and have a comprehensive system model of individual and collective device behaviors and dynamics, potentially limiting scalability, especially as the number of devices in the grid increases. The future of grid management increasingly leans toward distributed methods, where devices autonomously manage their energy consumption while coordinating with other elements of the grid. It is worthwhile to further investigate the benefits and trade-offs offered by distributed device-driven approaches, which could operate without complex models and have been shown to scale from literature to real-world applications [86, 87]. This dissertation, as we shall see in later chapters, focuses on advancing this shift by

developing practical, scalable, device-driven strategies that can operate effectively without disrupting user comfort, building on existing models like PEM.

There are significant efforts underway to decarbonize residential buildings by replacing fossil fuel-dependent thermostatic loads with electric alternatives like heat pumps and electric water heaters [88–90]. While these resulting electric loads present an increase in demand-side resources with ancillary service potential, they also have impacts on decarbonization, economic considerations and peak energy demand. Precisely and accurately modeling the energy dynamics of these electric loads are important in ensuring inaccurate model estimates do not undermine decarbonization assessments and policies. The modeling approaches used in harnessing the ancillary service potential, as in [56], are not suitable for these longer-term assessments due to significant differences in timescales—seconds or minutes versus months or years. Despite the extensive literature on energy modeling [91–95], there remains limited or no comprehensive review focused specifically on thermostatic electric loads in residential settings in the United States. This dissertation provides a comprehensive overview of existing energy methods and models used in understanding and quantifying the decarbonization and other non-ancillary benefits of electrification, highlighting their applicability and limitations.

1.3 Open Problems

The academic literature has extensively explored various load control strategies, including price-based incentives, direct load control, and smart home technologies. However, the transition from theoretical load control strategies to real-world implementation still faces many challenges not sufficiently addressed, such as distribution network issues, stability concerns and communication network issues that arise from coordinating distributed load resources at scale [55]. Addressing these will increase trust in load control technologies and the credibility of distributed demand-side load resources to provide balancing services to the power system. Distributed Energy Resource Management Systems (DERMS) are used in practice to facilitate the practical integration of diverse distributed energy resources, including load control, into

a unified and intelligent system. Commercially available DERMS provide functionalities such as real-time monitoring, predictive analytics, and automated control. Unfortunately, many of the control design in the academic literature are not easily adoptable in real-world DERMS due to practical constraints such as, limited accessibility to thermostat application programming interfaces, restricting data and state information available, restricted direct control in switching load states and restricted communication frequency. Reference [56] discusses these in detail. DERMS play a pivotal role in translating theoretical advancements into actionable solutions for grid operators and utilities. There is a need to overcome the practical barriers associated with DER control and also for control designs that can practically scale from the academic literature into real world DERMS. The synergy between renewable energy integration and VPPs via advanced load control strategies, underscores the importance of continued research, development, and implementation to address the evolving challenges in the energy landscape. As power systems transition towards a cleaner and more resilient future, the role of DSM through load control could be pivotal in shaping the trajectory of the energy transition. Hence, our focus in this dissertation is on practical and scalable load control and coordination strategies. While the academic literature has also extensively reviewed varying energy modeling approaches, there remains a gap with systematic reviews specifically concerning residential thermostatic electric loads, which is presently pertinent given the residential electrification and decarbonization initiatives [96].

Based on the discussion above, the following gaps in the academic literature on thermostatic load modeling and control, which this dissertation addresses, are:

1. Device-driven bottom-up load control strategies have not received enough attention in the literature. As the extensive literature review of Section 1.2 shows, aggregator-driven top-down load control strategies are often a popular choice, with perceived benefits of central visibility and controllability. However, these strategies face challenges of their own, as discussed in Section 1.2 above. It is worth examining device-driven strategies more rigorously, finding ways to leverage the practical benefits associated with their reduced modeling and state information requirements, and provide comprehensive analysis on the

trade-off between both aggregator-driven and device-driven approaches.

2. Sufficient analytical tools and models for robust device-driven bottom-up load control strategies are needed. Some of the control strategies, such as the packetized energy management for compressor-based loads [30], which is of interest to us in this work, does not have an analytical model. It is necessary to develop analytical approaches that can provide insights into the system dynamics and behavior, for a better understanding of the system beyond time-expensive simulations and heuristic predictions.

3. Practical demonstrations via hardware-in-the-loop and cyber-physical experiments are needed to better understand the controller dynamics in real-world settings, as well as demonstrate viability and scalability, which is vital to building trust in the ability of load control technologies to reliably offer balancing services such as frequency regulation.

4. A systematic analysis of energy modeling methods focused on flexible thermostatic loads is needed. Existing reviews either do not delve into the technical features and characteristics relevant to thermostatic load modeling or are too broad to provide practical guidance for researchers and practitioners working in this niche area. Previous research has predominantly concentrated on system-level energy-economic models or on very specific modeling frameworks.

1.4 Organization and Contributions

This dissertation contains six chapters. This current chapter - Chapter 1, serves as an introductory and background chapter. The summary of contributions and organization of the remainder of this dissertation is as follows:

Chapter 2 builds on the established device-driven control design in the literature called Packetized Energy Management (PEM). It improves the controller's capabilities, flexibility and performance whilst extending the controller compatible devices to include the class of thermostatic loads known as compressor-based TCLs. This

is an important subcategory of TCLs as more than 75% of residential TCL energy consumption is by compressor-based TCLs, and we know TCLs account for over 50% of residential energy use [38]. The controller modifications and extensions further enhance the practical viability of the control design. This work is completed and published in [30]. It addresses the first literature gap identified in Section 1.3.

In Chapter 3, we present a new practical bottom-up device-driven load control scheme with even better performance and flexibility than the state-of-the-art PEM controller and previous works discussed in Chapter 2. We also develop models describing the behavior of an aggregation of TCLs under the new control scheme and the state-of-the-art PEM; we discuss how this compares to the existing literature and the use cases of the model in understanding system dynamics, improving performance, predicting steady state conditions and required initialization conditions. Chapter 3 addresses the first literature gap in Section 1.3 by building on the existing work in the device-driven load control literature to develop a new practical bottom-up load control and coordination strategy. Also Chapter 3 addresses the second literature gap in Section 1.3 by creating analytical tools for better understanding device-driven control. This work is completed and the manuscript shall be imminently submitted.

Chapter 4 describes the hardware-in-the-loop experimental demonstration of load controllers with an experimental testbed consisting of 20 window-box air conditioners within physical model houses including a heat source and fan, and an associated communication network enabling control/coordination. This chapter discusses the design, development and validation of the load control experimental testbed, as well as present results from testing three load control methods - the PEM based control developed in Chapter 2, proportional integral control and a Markov-model based controller. In addition, it highlights the opportunities and challenges associated with testing TCL control strategies through physical experiments. It addresses the third literature gap identified in Section 1.3. This work is completed and under peer review [97].

In Chapter 5, we examine many modeling and simulation approaches for estimating energy dynamics of residential electric loads. It provides insights into data requirements, trade-offs, pros and cons of the different approaches, as well as the

suitable use-case for each energy modeling method. The discussion and analysis provides guidance to energy stakeholders on appropriate analytical and simulation tools based on project objectives and data availability. This is needed as models and software needed to understand real-time load control differs significantly from those for understanding decarbonization, economic and environmental impacts, particularly on larger time scales. It addresses the fourth literature gap identified in Section 1.3. This work is completed and the manuscript shall be imminently submitted.

Chapter 6 presents the conclusions of the dissertation and discusses future research directions.

In the paper [56] I co-authored under the leadership of my collaborator, Dr. Granitsas, we present a real-world, at-scale demonstration of load control and coordination involving approximately 100 homes in Texas. This demonstration integrated virtual TCLs to create a cyber-physical system combining physical and simulated TCLs. My primary contribution was the design and implementation of the communication architecture used within the cyber-physical system, detailed in Section II of the paper. Additional technical details of the communication setup are provided in Appendix A. Furthermore, I played a significant role in developing the high-fidelity simulation subsystem described in Section IV of the paper, which was pivotal to the cyber-physical demonstration. Specifically, I designed and implemented fully configurable communication networks modeled with configurable links, package drop rates and delays. This communication network model was then integrated with TCL and distribution network models, enabling co-simulation of TCLs on distribution networks coordinated over communication networks. This high-fidelity simulation subsystem is discussed in detail in Section IV of [56] and is publicly available at [98]. This work addresses the third literature gap identified in Section 1.3.

CHAPTER 2

Control of Aggregate Air-Conditioning Load using Packetized Energy Concepts

This chapter is based on work done and published in [30]. This chapter presents the extension of the packetized energy management (PEM) control strategy to enable coordination of compressor-based thermostatically controlled loads (TCLs), such as air conditioners. This establishes a new method of harnessing the flexibility of this ubiquitous resource, enabling a variety of grid services, such as frequency regulation. In the original PEM scheme, resources request energy packets and turn on if their request is approved. That PEM scheme has been further extended by introducing the concept of turn-off requests. We find that this increases flexibility and improves tracking performance. Through a case study involving over 1000 air conditioners, we evaluate the performance of a population of TCLs providing frequency regulation under PEM, highlighting both the capabilities and limitations. Simulations indicate our controller extensions significantly increase resource availability and tracking performance. We show that it is possible to achieve RMS tracking error below 2% when providing more than 250 kW of frequency regulation.

2.1 Notation

C_{a}	Air thermal mass [kWh/°C].
C_{m}	Building thermal mass [kWh/°C].
H_{m}	Conductance between inner air & solid mass [kW/(°C)].
k	Time step [-].
$\mathcal{K}_{\mathrm{lock}}$	Set of lock-out time-steps [-].
m	ON/OFF switching variable [-].
M_{on}	Mean on-request frequency [Hz].
M_{off}	Mean off-request frequency [Hz].
n	Air Conditioner index [-].
Q_{a}	Heat flux into interior mass [kW].
Q_{m}	Heat flux into interior solid mass [kW].
R_n^{on}	Energy packet request [-].
R_n^{off}	Turn-off request [-].
t^{on}	Epoch [s].
$t_{\mathrm{locked}}^{\mathrm{on}}$	Compressor turn-on lock-out time [s].
$t_{\mathrm{locked}}^{\mathrm{off}}$	Compressor turn-off lock-out time [s].
$t_{\mathrm{min}}^{\mathrm{on}}$	Energy packet minimum epoch length [s].
$t_{\mathrm{max}}^{\mathrm{on}}$	Energy packet maximum epoch length [s].
t_{comp}	Compressor lock-out timer [s].
t_n	Elapsed epoch time for AC n [s].
T_{a}	Indoor Air Temperature [°C].
T_{m}	Inner Mass Temperature [°C].
T_{o}	Outdoor Air Temperature [°C].

T_n^{set}	Temperature set-point [°C].
$T_n^{\min}$	Lower dead-band temperature [°C].
$T_n^{\max}$	Upper dead-band temperature [°C].
U_{a}	Conductance of building envelope [kW/°C].
Δt	Time step duration [s].
μ	Temperature based rate parameter [-].
γ	Time based rate parameter [-].

2.2 Chapter Introduction

Environmental concerns and the drive to decarbonize the power sector have resulted in traditional generation, such as coal fired power plants, being phased out and replaced by renewable energy alternatives. Continual growth in renewable generation requires an increase in the amount of balancing services needed to compensate for the inherently intermittent nature of renewable energy generation and maintain power grid reliability [99]. Traditionally, balancing services have been provided by synchronous generators, but new sources of flexibility are needed. Aggregations of distributed energy resources (DERs) such as thermostatically controlled loads (TCLs), and more specifically residential air conditioners (ACs), can offer such flexibility [2, 31, 64, 100–103].

The recent US Federal Energy Regulatory Commission Order 2222 [17] requires that grid operators revise their rules to enable DER aggregations to participate alongside other traditional resources in wholesale markets. This signals a paradigm shift in electricity markets that will lead to increased participation of DERs, such as flexible load aggregations. It ensures the need for load coordination and control strategies that meet both technical requirements and market participation criteria, and readily scale from the literature to real-world applications.

Houses have thermal inertia which allows ACs the flexibility to shift their energy consumption in time. Furthermore, ACs are ubiquitous resulting in a large resource potential; almost 90% of US homes utilize air conditioning equipment [42]. Usage of

residential ACs is also common in other countries [101, 104, 105]. The abundant and distributed geographic location of ACs make them prime and valuable candidates for providing flexibility to the grid, especially when aggregated.

Leveraging these AC resources for grid services requires coordination and control that is non-disruptive to occupants, robust, dynamic, and efficient, whilst satisfying concerns for privacy and autonomy. Numerous AC coordination and control schemes have been proposed. Some are centralized, others decentralized, some are model free, others model based, some are suitable for fast time scale dynamic signals, others can only work on slower time scales [2, 54, 79, 106–111]. However, there remains scope for improving effectiveness and practicality.

The goal of this chapter is to develop a packetized energy management (PEM) approach [61] to effectively control an aggregation of ACs to provide frequency regulation. PEM is a device-focused scheme that has been shown in a real-world application to achieve non-disruptive, asynchronous, and anonymous load coordination using electric resistance water heaters for load following, i.e., tracking a reference signal over an extended period (4 hours). However, the existing PEM scheme does not explicitly consider the capabilities and constraints of compressor-based TCLs like ACs, heat pumps, and refrigerators, which exhibit fundamentally different dynamics to resistance-based TCLs. Moreover, resistance-based TCLs (water and space heaters) are gradually being replaced by compressor-based TCLs because they are more efficient. Hence, there is a need to extend the PEM scheme to cater for compressor-based TCLs.

A significant amount of work has been undertaken on the PEM control scheme. A macro-model was developed in [62] for a homogeneous population of electric water heaters under PEM control. This macro-model was extended in [112] to a diverse group consisting of electric water heaters and energy storage systems operating under PEM. Analysis in [113] provided a discrete-time control law that maximizes requests accepted whilst tracking a regulation signal. This analysis is further extended in [114, 115] to accommodate diverse device types, including batteries, EVs, and resistance-based water heaters. Finally, [116] develops a decentralized frequency-responsive control scheme for electric water heaters, using local frequency-dependent control

policies based on an adaptation of PEM. However, none of this prior work has considered ACs with operational compressor constraints.

In this chapter, we develop a non-disruptive control approach to coordinate ACs under PEM, with the goal of allowing an aggregator to provide frequency regulation. Our approach accounts for compressor lock-out, i.e., compressors cannot be switched on immediately after switching off or switched off immediately after switching on. This restriction has been considered in other load control work, e.g., [117], but not PEM. We also improve capabilities of the PEM scheme by incorporating turn-off mechanisms, i.e., TCLs not only issue a request to turn on but can also request to turn off.

The chapter makes a number of contributions. First, we adapt the PEM control strategy to compressor-based cooling TCLs and introduce the idea of a flexible epoch (the duration of time for which the TCL is turned on). This mitigates the negative effect of lock-out on resource availability. Second, we extend the PEM control strategy to incorporate a mechanism for turning TCLs off mid-epoch, which increases flexibility and improves tracking performance. Third, we demonstrate PEM for compressor-based TCLs under realistic conditions, showing significant improvements in the fleet's performance score.

The rest of the chapter is structured as follows. Section 2.3 describes AC operation and our AC model, provides an overview of PEM, and proposes a modification to PEM that allows compressor-based loads, such as ACs, to be incorporated. Section 2.4 details the controller modifications needed to support turn-off requests. Case study results and analysis are presented in Section 2.5, while Section 2.6 provides concluding remarks.

2.3 Air Conditioner Control with PEM

2.3.1 AC Overview and Model

AC power consumption is determined by its ON/OFF state, which is governed by a thermostat. The thermostat utilizes a temperature dead-band $[T^{\min}, T^{\max}]$ to

achieve hysteretic control that regulates temperature around a set-point T^{set}. The dead-band confines the house temperature to a narrow comfort region around the set-point. For fixed power ACs, i.e., ACs that consume rated power when cooling and minimal power when not, a dead-band is necessary to prevent the AC from frequent cycling, which would lead to wear-and-tear and shorten its lifespan. The AC turns on (cools) when the temperature rises to $T^{\max}$, remains on until the temperature falls to $T^{\min}$, then switches off until the temperature rises again to $T^{\max}$. This cycle gives rise to alternating periods of rated power draw and little/no power draw. Fixed power ACs are dominant in the US.

We model ACs using the equivalent thermal parameter (ETP) model of [44–46, 118–120], with indoor air temperature T_{a} and inner mass temperature T_{m} dynamics given by,

$$\begin{aligned}\dot{T}_{\text{a}}(t) &= \frac{1}{C_{\text{a}}}(T_{\text{m}}(t)H_{\text{m}} - (U_{\text{a}} + H_{\text{m}})T_{\text{a}}(t) + Q_{\text{a}}(m(t)) \\ &\qquad + T_{\text{o}}(t)U_{\text{a}}), \\ \dot{T}_{\text{m}}(t) &= \frac{1}{C_{\text{m}}}\left(H_{\text{m}}(T_{\text{a}}(t) - T_{\text{m}}(t)) + Q_{\text{m}}\right),\end{aligned}$$

where C_{a} is the air thermal mass, H_{m} is the conductance between inner air and solid mass, U_{a} is the conductance of the building envelope, T_{o} is the outdoor air temperature, C_{m} is the building thermal mass, and Q_{m} is the heat flux to the interior solid mass. Function Q_{a} is the heat flux into the interior mass, which includes the internal and ambient heat gains and cooling from the AC, which in turn is a function of the AC's discrete on/off mode m. These dynamics can be expressed as a linear time-invariant system

$$\dot{x}(t) = Ax(t) + Bu(t),$$

where

$$x(t) = \begin{bmatrix} T_{\mathrm{a}}(t) & T_{\mathrm{m}}(t) \end{bmatrix}^T,$$
$$u(t) = \begin{bmatrix} Q_{\mathrm{a}}(m(t)) & T_{\mathrm{o}}(t) & Q_{\mathrm{m}} \end{bmatrix}^T,$$
$$A = \begin{bmatrix} -(U_{\mathrm{a}} + H_{\mathrm{m}})/C_{\mathrm{a}} & H_{\mathrm{m}}/C_{\mathrm{a}} \\ H_{\mathrm{m}}/C_{\mathrm{m}} & -H_{\mathrm{m}}/C_{\mathrm{m}} \end{bmatrix},$$
$$B = \begin{bmatrix} 1/C_{\mathrm{a}} & U_{\mathrm{a}}/C_{\mathrm{a}} & 0 \\ 0 & 0 & 1/C_{\mathrm{m}} \end{bmatrix}.$$

Using a timestep of τ, the above continuous-time state space model can be discretized as,

$$x[k+1] = A_d x[k] + B_d u[k],$$

where $A_d = e^{A\tau}$ and $B_d = A^{-1}(A_d - I)B$.

In discrete time, the on/off switching dynamics are,

$$m[k+1] = \begin{cases} 1, & \text{if } m[k] = 0 \ \wedge \ T_{\mathrm{a}}[k] \geq T^{\max} \\ 0, & \text{if } m[k] = 1 \ \wedge \ T_{\mathrm{a}}[k] \leq T^{\min} \\ m[k], & \text{otherwise.} \end{cases}$$

AC manufacturers also include compressor lock-out restrictions. After an AC has turned on it cannot turn off for a duration of $t^{\mathrm{on}}_{\mathrm{locked}}$ and after it has turned off it cannot turn on for a duration of $t^{\mathrm{off}}_{\mathrm{locked}}$. The lock-out times prevent the compressor from rapidly oscillating between states, protecting it from short cycling, and prolonging its lifespan.

2.3.2 Packetized Energy Management Overview

In PEM, TCLs make a request to turn on for fixed-duration periods to consume a fixed amount of energy. This has the effect of splitting and shifting their consumption in time such that they are on for shorter durations and switched more often. The concept is motivated by digital communications where large data packets are broken and transmitted in smaller packets. The fixed-duration energy consump-

tion is referred to as an *energy packet*. The fixed duration is called the *epoch*, t^{on}. TCLs choose whether to request an energy packet based on their air temperature and on/off state. The load aggregator coordinating the TCLs approves or denies the energy packet requests to achieve some overarching objective, such as using the aggregate power consumption of the TCLs to track a reference signal [61, 114, 115].

Fig. 2.1 shows the PEM scheme for TCLs. In this diagram, we see that the *Packetized Energy Controller* at the device in the consumer's home contains *Request Logic* that determines when the TCL should request an energy packet. This request logic utilizes local state information from TCLs to decide whether or not to make an *Energy Packet Request*. The aggregator processes the requests along with the *Desired Aggregate Power* of the fleet and uses its *Request Processing Logic* to decide whether to approve or deny each packet request. Since energy requests are initiated by the TCLs, they arrive stochastically at the aggregator. The requests are anonymous, i.e., the identity of requesting TCLs is encrypted and unavailable to the aggregator, thus preserving consumer privacy. The aggregator accepts/denies requests based on system needs. This overall process ensures fair resource deployment.

TCLs participating in PEM make energy requests automatically and probabilistically. Specifically, we adopt the approach from [61], where the probability that AC n makes an energy packet request R_n^{on} during the kth time-step (of duration Δt) is driven by the cumulative exponential distribution function,

$$P(R_n^{\text{on}}|T_{\text{a},n}[k]) := 1 - e^{-\mu(T_{\text{a},n}[k])\Delta t}, \tag{2.1}$$

where $T_{\text{a},n}$ is the indoor air temperature of the house conditioned by AC n and $\mu(T_{\text{a},n}[k])$ is a rate parameter based on the indoor air temperature (and defined below). We also take into account the AC's temperature dead-band to force an energy packet request if the house temperature is at the upper dead-band limit and ensure there is no request if the temperature is at the lower dead-band limit, i.e.,

$$P(R_n^{\text{on}} \mid T_{\text{a},n}[k] \leq T_n^{\min}) = 0, \tag{2.2}$$

$$P(R_n^{\text{on}} \mid T_{\text{a},n}[k] \geq T_n^{\max}) = 1. \tag{2.3}$$

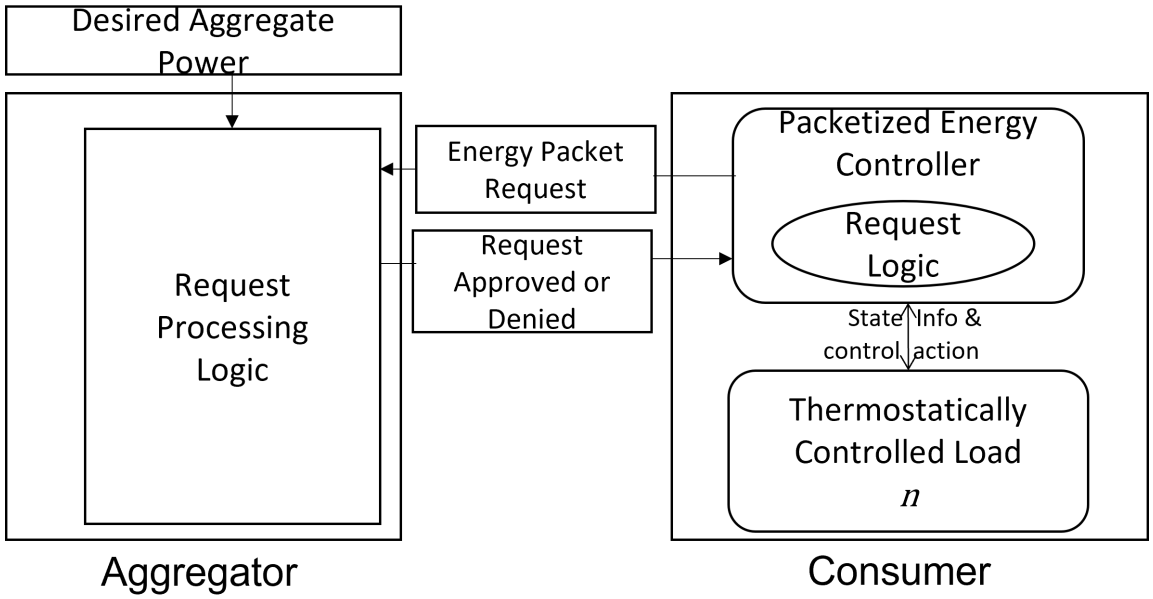


Figure 2.1: PEM Scheme for TCLs.

However, to ensure consumer comfort is always satisfied, the TCL reverts to default thermostat control when outside the dead-band, i.e., when $T_{\mathrm{a},n}[k] > T_n^{\max}$ the AC does not make an energy packet request but simply turns on.

The rate parameter is defined as,

$$\mu(T_{\mathrm{a},n}[k]) = \begin{cases} 0, & \text{if } T_{\mathrm{a},n}[k] \leq T_n^{\min} \\ f(T_{\mathrm{a},n}), & \text{if } T_{\mathrm{a},n}[k] \in (T_n^{\min}, T_n^{\max}) \\ \infty, & \text{if } T_{\mathrm{a},n}[k] \geq T_n^{\max} \end{cases} \tag{2.4}$$

where

$$f(T_{\mathrm{a},n}[k]) = \left(\frac{T_{\mathrm{a},n}[k] - T_n^{\min}}{T_n^{\max} - T_{\mathrm{a},n}[k]} \right) M_{\mathrm{on}} \tag{2.5}$$

and M_{on} is the mean on-request frequency at $T_{\mathrm{a},n}[k] = T_n^{\mathrm{set}} := (T_n^{\max} + T_n^{\min})/2$, which is a design parameter. A mean time to on-request of 10 min implies $M_{\mathrm{on}} = \frac{1}{600}$ Hz [61]. We slightly constrict the dead-band limits used within (2.2)–(2.5), creating

a PEM dead-band, to help prevent violation of the true TCL dead-band constraints. This is particularly important for TCLs that change status close to the dead-band limits and that are required to remain in their new status for a fixed duration [61]. The choice of constriction is a design choice based on the thermal dynamics and epoch duration.

2.3.3 Air Conditioners as a Packetized Load

The PEM scheme presented in Section 2.3.2 does not address compressor lock-out. We next describe a modification to PEM that caters for compressor-based loads. The controller is equipped with a timer t_{comp}, which turns on whenever the compressor changes status, and resets to zero at the conclusion of the relevant lock-out period. Specifically, over the lock-out period following AC turn on, $t_{\text{comp}} \in (0, t^{\text{on}}_{\text{locked}}]$, with t_{comp} resetting after $t^{\text{on}}_{\text{locked}}$. Likewise, over the lock-out period following AC turn off, $t_{\text{comp}} \in (0, t^{\text{off}}_{\text{locked}}]$ and resets after $t^{\text{off}}_{\text{locked}}$. We refer to the set of timesteps during which the AC is off and locked as $\mathcal{K}^{\text{off}}_{\text{lock}}$ and those when the AC is on and locked as $\mathcal{K}^{\text{on}}_{\text{lock}}$. The AC makes no PEM requests when locked, i.e., (2.1) is overridden by $P(R^{\text{on}}_n | T_{\text{a},n}[k]) := 0$ when $t_{\text{comp}} > 0$, and (2.2)-(2.3) are replaced by,

$$P\big(R^{\text{on}}_n | T_{\text{a},n}[k] \leq T^{\min}_n \vee k \in \mathcal{K}^{\text{off}}_{\text{lock}} \vee k \in \mathcal{K}^{\text{on}}_{\text{lock}}\big) = 0, \tag{2.6}$$

$$P\big(R^{\text{on}}_n | T_{\text{a},n}[k] \geq T^{\max}_n \wedge k \notin \mathcal{K}^{\text{on}}_{\text{lock}} \wedge k \notin \mathcal{K}^{\text{off}}_{\text{lock}}\big) = 1. \tag{2.7}$$

Note that the epoch must be chosen to be larger than the on lock-out time, $t^{\text{on}} > t^{\text{on}}_{\text{locked}}$. Otherwise, at the end of an epoch when the TCL was due to turn off, it would be prevented from doing so by the lock-out override, and would revert to default thermostat control.

ACs that are able to make PEM requests, i.e., operating within the PEM dead-band and not locked, are considered available. Compressor lock-out significantly reduces TCL availability as we will demonstrate through case studies in Section 2.5.2. Reduced availability impacts control performance, which motivates the controller extensions in the following section.

2.4 Controller Extensions

In the PEM scheme presented in the previous section, the aggregator has the ability to turn on TCLs (by approving energy packet requests) but cannot turn off TCLs should the need arise. This makes it hard to closely track downward trajectories of the reference signal, as the aggregator must rely solely on rejecting all new energy packet requests and waiting for the unforced reduction in aggregate power consumption as AC epochs expire. This issue is compounded by compressor lock-out, which reduces TCL availability since TCLs are locked and unavailable as soon as they switch off at the end of an epoch. Reduced availability makes it harder to track in general. We will illustrate the impact of these limitations through case studies in Section 2.5.2.

To mitigate the effects of these practical challenges, we introduce a flexible epoch and a turn-off mechanism that increase flexibility and TCL availability. The epoch length t^{on} now lies in the range $[t^{\text{on}}_{\min}, t^{\text{on}}_{\max}]$ where $0 < t^{\text{on}}_{\text{locked}} \leq t^{\text{on}}_{\min} \leq t^{\text{on}}_{\max}$. The controller is equipped with a timer that records the elapsed epoch duration of an energy packet, $t \in (0, t^{\text{on}}_{\max}]$. An AC with an accepted energy packet request will turn on for an epoch length that is at least $t^{\text{on}}_{\min}$ and no greater than $t^{\text{on}}_{\max}$.

The turn-off mechanism enables turn-off requests R^{off}_n during the epoch range $(t^{\text{on}}_{\min}, t^{\text{on}}_{\max}]$. The probability that AC n in the ON state, operating in the PEM scheme, having an elapsed epoch time of $t_n \in (0,\ t^{\text{on}}_{\max}]$, makes a turn-off request R^{off}_n during the kth time-step is defined by,

$$P\big(R^{\text{off}}_n | T_{\text{a},n}[k], t_n[k]\big) := \big(1 - e^{-\mu^{\text{off}}(T_{\text{a},n}[k])\Delta t}\big)\big(1 - e^{-\gamma(t_n[k])\Delta t}\big), \tag{2.8}$$

where the turn-off temperature-based rate parameter is,

$$\mu^{\text{off}}(T_{\text{a},n}[k]) = \begin{cases} 0, & \text{if } T_{\text{a},n}[k] \geq T^{\max}_n \\ \left(\frac{T^{\max}_n - T_{\text{a},n}[k]}{T_{\text{a},n}[k] - T^{\min}_n}\right) M_{\text{off}}, & \text{if } T_{\text{a},n}[k] \in (T^{\min}_n, T^{\max}_n) \\ \infty, & \text{if } T_{\text{a},n}[k] \leq T^{\min}_n \end{cases} \tag{2.9}$$

the time-based rate parameter is,

$$\gamma(t_n[k]) = \begin{cases} 0, & \text{if } t_n[k] \leq t^{\text{on}}_{\min} \\ \left(\frac{t_n[k]-t^{\text{on}}_{\min}}{t^{\text{on}}_{\max}-t_n[k]}\right) M_{\text{off}}, & \text{if } t_n[k] \in (t^{\text{on}}_{\min}, t^{\text{on}}_{\max}) \\ \infty, & \text{if } t_n[k] \geq t^{\text{on}}_{\max} \end{cases} \tag{2.10}$$

and M_{off} is the mean off-request frequency, a design parameter similar to M_{on} described in Section II and detailed in [61]. Then the turn-off request is constrained as,

$$P\big(R^{\text{off}}_n | T_{\text{a},n}[k] \leq T^{\min}_n \wedge t_n[k] \geq t^{\text{on}}_{\max} \wedge k \notin \mathcal{K}^{\text{on}}_{\text{lock}} \wedge k \notin \mathcal{K}^{\text{off}}_{\text{lock}}\big) = 1, \tag{2.11}$$

$$P\big(R^{\text{off}}_n | T_{\text{a},n}[k] \geq T^{\max}_n \vee t_n[k] \leq t^{\text{on}}_{\min} \vee k \in \mathcal{K}^{\text{on}}_{\text{lock}} \vee k \in \mathcal{K}^{\text{off}}_{\text{lock}}\big) = 0. \tag{2.12}$$

To ensure consumer comfort, the TCL reverts to default thermostat control and does not make turn-off requests when outside the temperature dead-band, similar to Section 2.3.2. It also does not make turn-off requests when the AC is locked. The turn-off mechanism is consistent with the distributed nature of the PEM algorithm with request probabilities computed locally. Equation (2.9) ensures TCLs closer to the lower dead-band temperature have higher probabilities of making turn-off requests and (2.10) enables TCLs with higher elapsed epoch duration to have higher probabilities of making turn-off requests.

Fig. 2.2 shows the closed loop feedback system for a load aggregator with ACs under PEM control. The load aggregator receives a reference signal from the grid operator and must then coordinate the available resources to match that signal. Thus, at each sample period k, the load aggregator has two queues of anonymous energy packet requests (henceforth referred to as turn-on requests) and anonymous turn-off requests from ACs under its control. If the aggregate power consumption of the AC population under the load aggregator's control under- or over-shoots the reference signal, the aggregator first determines the number of turn-on or turn-off requests, respectively, to approve to minimize the tracking error. This is achieved by dividing the tracking error by an estimate of the average AC power consumption.

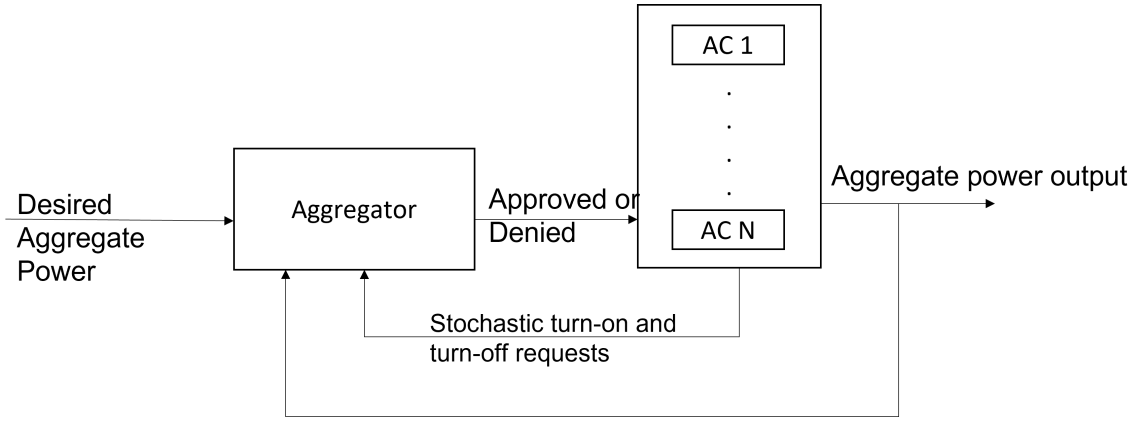


Figure 2.2: Extended PEM for closed-loop control of an AC aggregation.

The aggregator then selects and approves the desired number of ON/OFF requests from the corresponding queue. All other requests are denied.

The fixed epoch design of the original PEM scheme means that devices with approved packet requests are effectively unavailable as they cannot make further PEM requests during the epoch duration. The flexible epoch and turn-off mechanism allows these devices to become available for control once they have met their minimum epoch time. Because they can start making PEM turn-off requests, overall TCL availability is increased. This in turn improves the range of corrective actions (flexibility) and tracking performance as the aggregator has access to more requests. Also, the turn-off requests provide the aggregator with the ability to actively track downward reference trajectories. Specifically, in addition to rejecting new turn-on requests and the unforced reduction in aggregate power consumption due to epochs expiring, the aggregator can now actively turn off ACs mid epoch. The benefits of this new approach are demonstrated through the following case studies.

Table 2.1: Simulation Parameters

Parameters	Value	Parameters	Value
Δt	2 s	t^{on}	10 min
$t^{\text{on}}_{\text{locked}}$	3 min	$t^{\text{off}}_{\text{locked}}$	5 min
$t^{\text{on}}_{\text{min}}$	3 min	$t^{\text{on}}_{\text{max}}$	10 min
T^{set}_n	$[20, 24]^\circ$C	$T^{\text{max}}_n - T^{\text{min}}_n$	$[1, 2]^\circ$C
T_{o}	32.22°C	M_{off}	1 Hz
M_{on}	$\frac{1}{300}$ Hz	Simulation Period	1 h

2.5 Case Studies

2.5.1 Simulation Scenarios and Parameters

The case studies consider the response of 1103 ACs that are controlled using the PEM scheme to provide frequency regulation. The parameters of the ACs were generated using GridLAB-D [118], with $\pm 10\%$ random variation around each parameter. Each AC consumes 2.5 kW on average. This study considers houses with single central AC units, but is directly applicable and extendable to houses with multiple AC units given that every AC unit will make its own PEM request. The PJM RegD signal [121] was used as the reference signal to be tracked. Table 2.1 lists the simulation parameters that were used. As explained in Section 2.3.2, we constrict the dead-band used by the PEM controller, referred to as the PEM dead-band. Here we set the PEM dead-band width to $0.8(T^{\text{max}}_n - T^{\text{min}}_n)$ for all ACs, where 0.8 is a design choice that has been found to sufficiently mitigate violations of the AC thermostat dead-bands.

We assess the AC population's tracking performance under a variety of conditions. Specifically, we define seven cases with varying amounts of regulation capacity, with and without the controller extensions (flexible epoch and turn-off capability), and with and without compressor lock-out.

Table 2.2: Controller performance results across all cases

	Case 0	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
Regulation Capacity (MW)	0.25	0.25	0.25	0.5	0.5	1	1
Flexible Epoch and Turn-off Capability	No	No	Yes	No	Yes	No	Yes
Compressor Lock-out	No	Yes	Yes	Yes	Yes	Yes	Yes
NRMSE (%)	5.45	5.59	1.62	15.11	12.50	36.01	35.29
PJM Accuracy Score (0-1)	0.88	0.87	0.99	0.80	0.92	0.76	0.79
PJM Delay Score (0-1)	0.98	0.97	1.00	0.94	0.98	0.93	0.95
PJM Precision Score (0-1)	0.73	0.70	0.88	0.55	0.77	0.39	0.62
PJM Composite Score (0-1)	0.87	0.85	0.96	0.77	0.89	0.69	0.78
Mean Availability (%)	44.17	18.88	74.18	19.03	61.14	18.87	47.68

2.5.2 Simulation Results

Table 2.2 and Fig. 2.3 summarize the controller performance results across all of the cases. Table 2.2 gives the normalized root mean square error (NRMSE) between the aggregate power and the power reference, normalized by the baseline power consumption of the aggregation. We also provide an industry performance metric, namely PJM performance scores [122], and the mean availability, which is calculated as the average percentage of available devices over the simulation period. An AC is considered available when it is unlocked and within the PEM temperature deadband, and hence available to participate and make turn-on or turn-off requests under the PEM scheme. The top set of plots in Fig. 2.3 show the reference signal that is to be tracked, the aggregate power of the ACs, and the baseline power consumption. The bottom set of plots show the percentage of available ACs, the ramp-up flexibility, and the ramp-down flexibility. The ramp-up/down flexibility is the fraction of available ACs making turn-on/off requests. It defines the flexibility the aggregator has in increasing/decreasing aggregate power consumption by approving turn-on/off requests.

Fig. 2.4 shows the fraction of the population that is ON/OFF and the portion of those that are available/unavailable for Cases 0–2. Table 2.3 presents these values averaged over the simulation period. By comparing the results of Cases 0 and 1, we can see the impact of compressor lock-out on availability, which in turn affects performance. We see in Case 0 that an average of 55.8% ACs are ON while the remaining 44.2% are OFF and available because lock-out is ignored in Case 0. In Case 1, an average of 54.3% of ACs are ON, but of the 45.7% that are OFF, only 18.9% are available because of compressor lock-out. Hence, lock-out reduces the mean AC availability by a factor of 134%. This disparity in availability accounts for the noticeable difference in the ramp-up flexibility between both cases (shown in Fig. 2.3) as there are fewer available ACs making requests. Due to the reduced flexibility, we see from Table 2.2 that the PJM composite score decreases from 0.87 to 0.85 and the NRMSE increases from 5.45% to 5.59%. This effect becomes more pronounced when regulation capacity increases.



Figure 2.3: Reference tracking, resource availability, and flexibility for all cases.

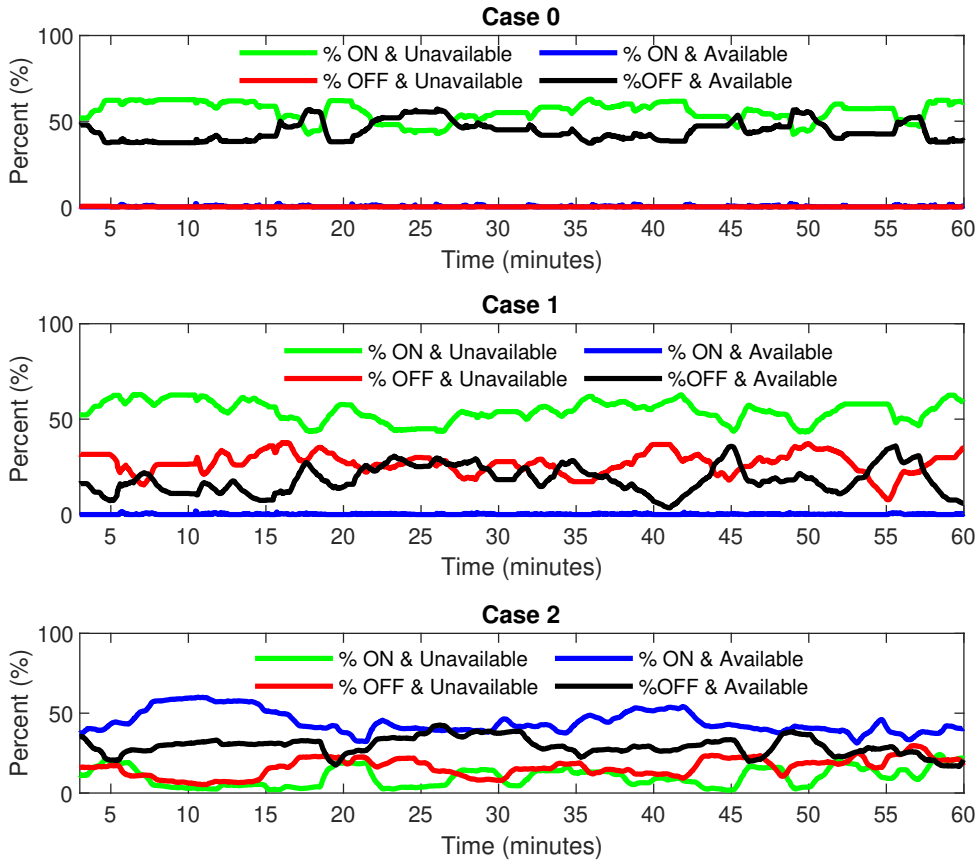


Figure 2.4: Effect of lock-out and controller extensions on AC availability.

Table 2.3: Average AC availability

	Case 0	Case 1	Case 2
ACs ON (%)	55.8	54.3	54.2
ACs OFF (%)	44.2	45.7	45.8
ON & Unavailable (%)	55.8	54.3	10.0
ON & Available (%)	0	0	44.2
OFF & Unavailable (%)	0	26.8	15.9
OFF & Available (%)	44.2	18.9	30.0

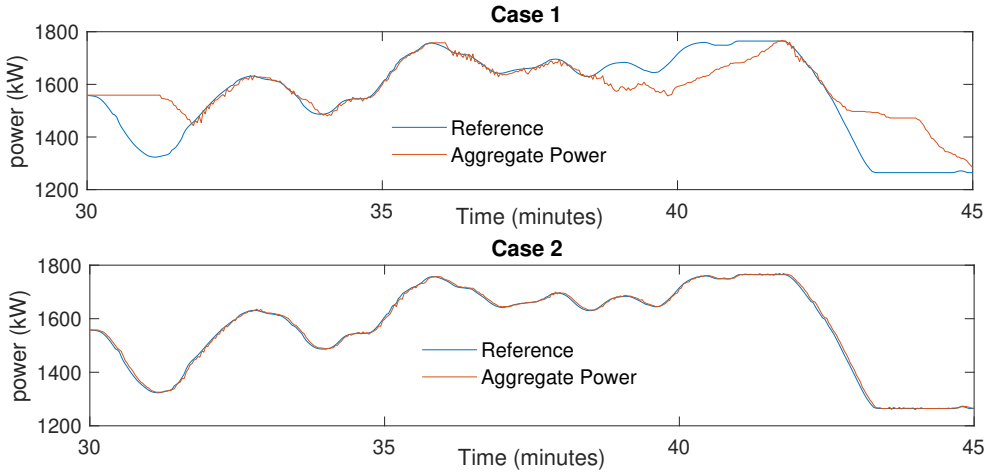


Figure 2.5: Close-up of the tracking results for Cases 1 and 2.

By comparing the results of Cases 1 and 2, we see that the extended controller mitigates the influence of lock-out on availability. We see from Table 2.3 that this results in an improvement of the mean device availability by a factor of 293%, which in turn affects the ramp-up/down flexibility, as seen in Fig. 2.3. This improved flexibility yields an increase in PJM composite performance score from 0.85 to 0.96 and a decrease in NRMSE from 5.59% to 1.62%. Fig. 2.5 highlights a 15 min window of the tracking results shown in Fig. 2.3 for Cases 1 and 2, clearly showing the impact of the controller extension on tracking performance.

Comparing the results of Cases 1–2, Cases 3–4, and Cases 5–6, we see that the controller extensions greatly increase device availability and flexibility. In Fig. 2.3, we see instances where the standard PEM controller is unable to track an upward trajectory as it is limited by the available ACs and ramp-up flexibility, and also unable to actively track downward trajectories without the turn-off mechanism. For example, comparing Cases 3 and 4, we see from Table 2.2 that the NRMSE decreases from 15.11% to 12.50%, mean device availability increases by a factor of 221.3%, and the PJM composite score increases from 0.77 to 0.89. Comparing Cases 5 and 6, we see a decrease in NRMSE from 36.01% to 35.29%, an increase in mean device availability by a factor of 152.7%, and an increase in the PJM composite score from

0.69 to 0.78.

PJM requires composite scores greater than 0.75 [122] and we see that the extended PEM controller qualifies for market participation with 0.25, 0.5, or 1 MW of regulation capacity. The original controller falls short in the 1 MW scenario (Case 5) because of the limited number of ACs. From Fig. 2.3 we see the tracking is often hampered by the available ramp-up and ramp-down capacity. More devices are simply needed.

To improve the performance in the 1 MW frequency regulation capacity case, we adjusted M_{on} to increase the turn-on request rate while holding all other parameters constant. This increases the number of ACs available to the aggregator. Fig. 2.6 shows the new tracking performance; the PJM composite score increases from 0.78 (Case 6 in Table 2.2) to 0.84. Performance is limited and constrained by the number of available ACs. Accordingly, it is important to be able to ascertain the number of ACs needed to offer a given regulation capacity or the maximum capacity attainable by a given population of ACs, along with the corresponding optimal parameters, whilst taking into account the required performance objectives. Estimating the flexibility and choosing the optimal parameters heuristically via multiple simulations is cumbersome and inefficient due to the large parameter space [123]. Hence, developing an aggregate model that accurately approximates the dynamics of compressor-based TCLs is subsequently explored in Chapter 3. The flexibility and optimal parameter selection is impacted by ambient conditions. For example, the flexibility under nominal conditions will differ significantly from that under extreme heat conditions when most ACs will be constantly ON. Performance, flexibility, and optimal parameter selection under extreme conditions is explored in Chapter 3.

2.6 Conclusion

In this chapter, we extended the PEM aggregate control scheme to incorporate compressor-based TCLs. We included compressor lock-out constraints and developed both flexible epochs and a turn-off mechanism. We showed that these extensions improve device availability, flexibility, and controller performance. We found that

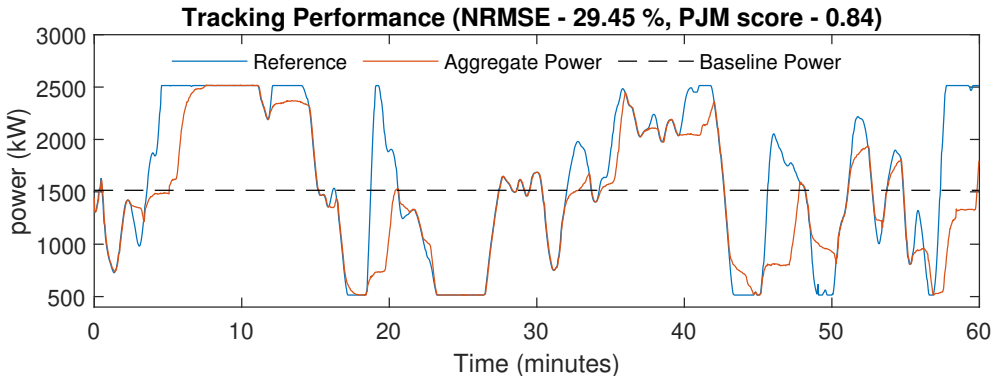


Figure 2.6: Extended PEM providing 1 MW frequency regulation, with new M_{on}. PJM composite score is 0.84.

with a population of around 1000 ACs, it is possible to offer up to 1 MW of frequency regulation capacity while meeting PJM market qualifying scores.

An aggregate model that accurately approximates the dynamics of compressor-based TCLs under PEM is yet to be developed. The existing macro-models described in the introduction do not account for compressor-based loads like ACs under PEM control. Such a model is important to analytically ascertain the flexibility limits of a fleet of compressor-based TCLs under PEM. An aggregate model is also vital for analytically determining the population steady-state distribution and analyzing dynamic behavior. Additionally, it would allow sensitivity analysis and optimal design of parameters such as epoch and mean request frequencies.

The PEM scheme requires bi-directional communication links between devices and the aggregator compared to some other schemes that employ unidirectional message broadcasting. Subsequent chapters explore controller characteristics and performance in the presence of communication latency and drop-outs, and will develop strategies to mitigate issues associated with these real-world challenges. We will also consider ways of enhancing the PEM controller to achieve distribution-network awareness. Furthermore, we will benchmark PEM against other state-of-the-art control strategies to understand trade-offs between centralized methods and more distributed schemes like model-free PEM.

CHAPTER 3

Device-Driven Modeling and Control of Thermostatic Loads: A Bottom-Up Approach

Actively managing loads in modern power systems could provide intermittent renewable generation with the balancing support it needs. Flexible loads could be pivotal in maintaining power system reliability and unlocking the potential for increased renewables integration. This chapter builds upon the demand dispatch paradigm known as Packetized Energy Management (PEM). PEM utilizes an energy packet-based concept to coordinate demand-side resources for power systems operational objectives. This work creates a new and more flexible device-driven load control strategy by eliminating the heuristic time-based requirements of PEM. We comprehensively assess our device-driven strategy applied to thermostatically controlled loads (specifically, air conditioners) under testing conditions ranging from ideal to extreme, and benchmark against the state-of-the-art PEM controller. Simulations show that tracking error, resource availability, and flexibility improve significantly with the new control strategy. A dynamical system model of an aggregation of loads operating under the device-driven strategy that can be used for predictive control and analysis is also developed.

3.1 Chapter Introduction

Decarbonization of the power sector will phase out most traditional power generation such as coal-fired power plants and replace them with renewable energy alternatives like wind and solar. The inherent intermittent nature of these renewable energy sources means the traditional approaches of maintaining power system reliability (i.e., managing generation to match scheduled demand) must be augmented with new sources of flexibility. Virtual power plants that coordinate distributed energy resources like thermostatically controlled loads (TCLs) – including air conditioners (ACs), heat pumps, and water heaters – can offer this flexibility [31].

There is a large body of research on hierarchical TCL control for power system objectives, most of which can effectively be categorized as either "aggregator-driven" [2] or "device-driven" [124] based on the load control strategy. The aggregator-driven approach entails primary computations and decisions being taken by the load controlling entity called the aggregator, while in the device-driven approach, these computations/decisions primarily occur at the device (load) level. The most effective and practical approach to implementing load control within the device-driven paradigm remains an open problem, and is a focus of this work.

In this chapter, we develop a new device-driven approach inspired by Packetized Energy Management (PEM) [61,62], which uses the packet-based concept from digital communications to control and coordinate device energy consumption. Specifically, PEM uses a request-response mechanism that allows devices to probabilistically request energy based on individual device states and parameters, and the aggregator approves (or denies) the energy request for a fixed time duration referred to as an epoch, according to operational needs. In PEM, fixed-duration energy consumption of devices are "packetized" (broken into smaller time durations) so energy use can be flexibly shifted in time, enabling the aggregator to achieve an operational objective, such as tracking a power reference signal with the aggregate power consumption of controlled devices. However, the approach was not tailored to TCLs with compressors. Most modern TCLs are compressor-based, though non-compressor-based TCLs, such as resistive space and water heaters, are still used. Compressor-based TCLs re-

quire lock-out constraints, meaning compressors switching on/off must remain on/off for several minutes to avoid short-cycling and damage to the compressor [125]. The PEM scheme in [61, 62], which we subsequently refer to as the "basic PEM scheme," did not consider lock-out.

Our prior work [30] extended PEM to accommodate lock-out by introducing a flexible epoch. It also enhanced the request probability design to enable active downward signal tracking. However, the design relied on heuristic selection of parameters (like the epoch), requiring re-examination. Also, the request probability design has a significant impact on modeling complexity and performance, requiring further analysis. Furthermore, while [63, 112, 114] developed aggregate system models for PEM, these models cannot be used to capture the controlled dynamics resulting from the PEM scheme in [30] and so there is need for new aggregate system models that help us characterize and analyze the dynamics of PEM extensions.

Therefore, here, we address these gaps through the design of a new device-driven scheme. Specifically, we strip the PEM controller of its inherent packetized approach, which eliminates the need for heuristic parameter selections. The new controller is applicable to both compressor-based and non-compressor based TCLs. We consider the impact of controller request probability design on performance and modeling complexity, and demonstrate that our controller is simpler, has better tracking performance, and offers increased flexibility when compared to the approach in [30] under a wide range of scenarios. Also, we develop a reduced-form aggregate system model for approximating the behavior of an ensemble of ACs under the new controller, and we show how this model can help us analyze and predict the dynamics of this energy resource and new device-driven controller. Additionally, we discuss how the model can be extended to capture the dynamics associated with the PEM scheme developed in [30].

In summary, our contributions are threefold: 1) we develop a new device-driven controller that has better performance and greater flexibility than the state-of-the-art PEM controller, 2) we develop a new aggregate system model of the controlled system dynamics, and 3) we present simulation results showing the benefits of the model and control scheme under ideal and extreme test conditions.

The chapter is structured as follows. Section 3.2 describes the control framework. Section 3.3 discusses the shortcomings of the previous controller and develops the new controller. Section 3.4 develops an aggregate model of the dynamics of the controlled system. Case study results and analysis are presented in Section 3.5, while Section 3.6 provides concluding remarks.

3.2 Framework

In aggregator-driven load control strategies, the aggregator uses available data, models, and state information to estimate the flexibility of the ensemble of devices it coordinates. The aggregator receives commands such as tracking signals from the system operator, and then formulates a control action based on the operator's command and modeled flexibility. Control inputs are sent in a "top-down" manner from the aggregator to all devices. Control inputs may be individual (tailored to each device) or general (the same to all devices). This control strategy is termed aggregator-driven because most computations and decisions take place at the aggregator level.

Alternatively, in device-driven load control strategies, most computations and decisions are made at the device level. Devices determine locally (based on their own state) whether or not to offer their service to (e.g., request energy from) the aggregator. Decisions are fed upwards in a "bottom-up" manner to the aggregator. Consequently, the aggregator can ascertain the flexibility of the ensemble of devices it coordinates. The aggregator then decides which devices to engage based on the system operator's signal. Examples of device-driven control include PEM [30, 61, 62] and earlier packetized direct load control approaches [126, 127]. Note that device-driven control can also be done without central aggregators, as in [128, 129]. In these decentralized approaches, the devices operate according to local sensing and control objectives only. However, as noted in [124], these methods are often constrained in their capability compared to the aggregator-involved device-driven control and aggregator-driven control, owing to limited system level information.

Here, we focus on device-driven control, which has been shown to scale to real-

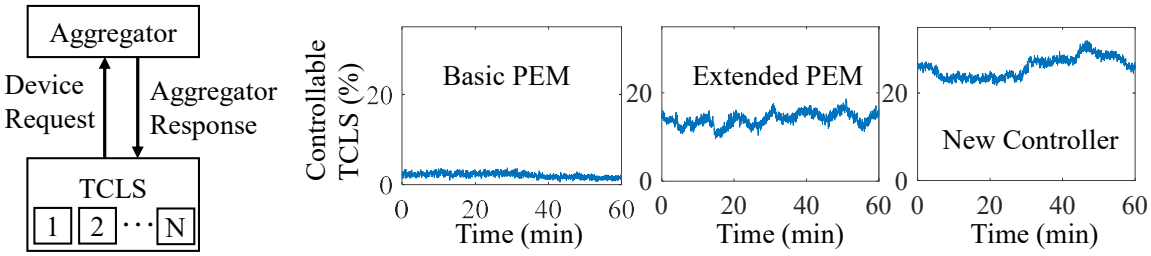


Figure 3.1: Device-driven load control architecture and flexibility across the different strategies.

Table 3.1: System Architectures across device-driven load controllers.

	Device Request	Aggregator Response
Basic PEM	Turn-On: temperature-based	Turn-On approvals: time-constrained
Extended PEM	Turn-On: temperature-based Turn-Off: temperature & time-based	Turn-On approvals: time-constrained Turn-Off approvals: unconstrained
New controller	Turn-On: temperature-based Turn-Off: temperature-based	Turn-On approvals: unconstrained Turn-Off approvals: unconstrained

world applications [86, 87]. Also, we focus on the compressor-based class of TCLs, an important subcategory of TCLs, as they account for more than 75% of the energy consumed by residential TCLs, and TCLs account for over 50% of residential energy use in the US [38]. In particular, for this study, we will consider air conditioners (ACs) though our methods are applicable to all TCLs.

Fig. 3.1 and Table 3.1 illustrate our proposed control system architecture and how it compares to two other forms of PEM, basic and extended PEM. All device requests within the PEM schemes are based on device local state information such as temperature or the length of time the device has spent within its current state, as will be detailed in Section 3.3.

In the basic PEM control scheme [61, 62], devices make turn-on (energy consumption) requests to the aggregator, who either approves or denies these requests depending on the overall objective. Approved turn-on requests are time-constrained

in that they enable a device to consume energy for a fixed duration called the epoch. Fig. 3.1 depicts a typical percentage of devices that are controllable under basic PEM, that is, the fraction of devices making turn-on requests at each interval over an illustrative one hour time horizon.

In the extended PEM control scheme [30], devices make both turn-on and turn-off (zero energy consumption) requests to the aggregator. The requests are either denied or approved. Here, the energy consumption approvals have a flexible epoch within a minimum and maximum duration. This flexibility allows turn-off requests to be made. Approved turn-off requests are unconstrained in that a device can stay off for an arbitrary amount of time. The controllable percentage of TCLs under the extended PEM scheme is shown in Fig. 3.1 for an illustrative one hour time horizon. As shown, the aggregator has access to more flexibility with extended PEM.

In both basic and extended PEM, epoch durations are heuristically chosen by the aggregator. Ref. [87] shows that estimating and choosing the optimal epoch is cumbersome and inefficient due to the large parameter space. Therefore, this chapter introduces a control design where devices still make turn-on or turn-off requests but the approvals are not time-constrained. The new controller keeps the request-response approach of PEM but does not employ the packetized concept. As shown in Fig. 3.1, the aggregator has access to even more flexibility with the new controller.

3.3 Device-Driven Load Control for Air Conditioners

This section details the modeling and control framework for ACs under the PEM paradigm and the proposed approach under the new controller.

3.3.1 Individual AC Model

ACs are modeled using the equivalent thermal parameter (ETP) model of [45, 118–120] with indoor air temperature T_{a} and inner mass temperature T_{m} dynamics

given by,

$$\begin{aligned}\dot{T}_{\mathrm{a}}(t) &= \frac{1}{C_{\mathrm{a}}}(T_{\mathrm{m}}(t)H_{\mathrm{m}} - (U_{\mathrm{a}} + H_{\mathrm{m}})T_{\mathrm{a}}(t) + Q_{\mathrm{a}}(m(t)) \\ &\qquad + T_{\mathrm{o}}(t)U_{\mathrm{a}}), \\ \dot{T}_{\mathrm{m}}(t) &= \frac{1}{C_{\mathrm{m}}}\left(H_{\mathrm{m}}(T_{\mathrm{a}}(t) - T_{\mathrm{m}}(t)) + Q_{\mathrm{m}}\right),\end{aligned}$$

where C_{a} is the air thermal mass, H_{m} is the conductance between inner air and solid mass, U_{a} is the conductance of the building envelope, T_{o} is the outdoor air temperature, C_{m} is the building thermal mass, and Q_{m} is the heat flux to the interior solid mass. Function Q_{a} is the heat flux into the interior mass, which includes the internal and ambient heat gains and cooling from the AC, which in turn is a function of the AC's discrete on/off mode m. These dynamics expressed as a linear time-invariant system:

$$\dot{x}(t) = Ax(t) + Bu(t),$$

where

$$\begin{aligned}x(t) &= \begin{bmatrix} T_{\mathrm{a}}(t) & T_{\mathrm{m}}(t) \end{bmatrix}^T, \\ u(t) &= \begin{bmatrix} Q_{\mathrm{a}}(m(t)) & T_{\mathrm{o}}(t) & Q_{\mathrm{m}} \end{bmatrix}^T, \\ A &= \begin{bmatrix} -(U_{\mathrm{a}} + H_{\mathrm{m}})/C_{\mathrm{a}} & H_{\mathrm{m}}/C_{\mathrm{a}} \\ H_{\mathrm{m}}/C_{\mathrm{m}} & -H_{\mathrm{m}}/C_{\mathrm{m}} \end{bmatrix}, \\ B &= \begin{bmatrix} 1/C_{\mathrm{a}} & U_{\mathrm{a}}/C_{\mathrm{a}} & 0 \\ 0 & 0 & 1/C_{\mathrm{m}} \end{bmatrix}.\end{aligned}$$

Using a time step of τ, the above continuous-time state space model can be discretized as,

$$x[k+1] = A_d x[k] + B_d u[k],$$

where $A_d = e^{A\tau}$ and $B_d = A^{-1}(A_d - I)B$.

In discrete time, the on/off switching dynamics are,

$$m[k+1] = \begin{cases} 1, & \text{if } m[k] = 0 \ \wedge \ T_{\mathrm{a}}[k] \geq T^{\max} \\ 0, & \text{if } m[k] = 1 \ \wedge \ T_{\mathrm{a}}[k] \leq T^{\min} \\ m[k], & \text{otherwise.} \end{cases}$$

The on/off state determines the AC power consumption, and the on/off dynamics are dictated by a thermostat. The thermostat utilizes a temperature dead-band $[T^{\min}, T^{\max}]$ to achieve hysteretic control that regulates temperature around a set-point T^{set}. The dead-band confines the house temperature to a narrow comfort region around the set-point. The AC turns on (cools) when the temperature rises to $T^{\max}$, remains on until the temperature falls to $T^{\min}$, then switches off until the temperature rises again to $T^{\max}$. This cycle gives rise to alternating periods of AC power consumption and no power consumption.

3.3.2 Extended PEM Overview and Shortcomings

In the extended PEM scheme, turn-on request probability $P\big(R_n^{\mathrm{on}}\big)$ at each time-step k (of duration Δt) is evaluated according to the cumulative exponential distribution function,

$$R_n^{\mathrm{on}}(T_{\mathrm{a},n}[k]) := 1 - e^{-\mu(T_{\mathrm{a},n}[k])\Delta t}, \tag{3.1}$$

where $T_{\mathrm{a},n}$ is the indoor air temperature being cooled by AC n and $\mu(T_{\mathrm{a},n}[k])$ is a temperature-based rate function, and is defined below,

$$\mu(T_{\mathrm{a},n}[k]) = \begin{cases} 0, & \text{if } T_{\mathrm{a},n}[k] \leq T_n^{\min} \\ f(T_{\mathrm{a},n}), & \text{if } T_{\mathrm{a},n}[k] \in (T_n^{\min}, T_n^{\max}) \\ \infty, & \text{if } T_{\mathrm{a},n}[k] \geq T_n^{\max} \end{cases} \tag{3.2}$$

where

$$f(T_{\mathrm{a},n}[k]) = \left(\frac{T_{\mathrm{a},n}[k] - T_n^{\min}}{T_n^{\max} - T_{\mathrm{a},n}[k]} \right) M_{\mathrm{on}} \tag{3.3}$$

and M_{on} is the mean on-request frequency, a design parameter affecting how often turn-on requests are made. A mean turn-on request frequency of 10 mins implies $M_{\text{on}} = \frac{1}{600}$ Hz.

To ensure consumer comfort is always satisfied, the AC reverts to default thermostat control when outside the dead-band, i.e., when $T_{\text{a},n}[k] > T_n^{\max}$ the AC does not make a turn-on request but simply turns on. We slightly constrict the dead-band limits used within (3.2) and (3.3), creating a PEM dead-band, to help prevent violation of the true AC dead-band constraints. This is particularly important for ACs close to the dead-band limits with successful requests (low probability but not zero), changing states and are required to remain in their new states for a fixed duration, per compressor lock-out constraints. The constriction width is a design choice based on the thermal dynamics and epoch duration. Taking into account the temperature dead-band limits of the AC and the compressor lock-out dynamics constrains the probability request function as follows,

$$P\big(R_n^{\text{on}} | T_{\text{a},n}[k] \leq T_n^{\min} \vee k \in \mathcal{K}_{\text{lock}}^{\text{off}} \vee k \in \mathcal{K}_{\text{lock}}^{\text{on}}\big) = 0, \tag{3.4}$$

$$P\big(R_n^{\text{on}} | T_{\text{a},n}[k] \geq T_n^{\max} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{on}} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{off}}\big) = 1. \tag{3.5}$$

where $\mathcal{K}_{\text{lock}}^{\text{on}}$ is the set of timesteps during which the AC is on and locked and $\mathcal{K}_{\text{lock}}^{\text{off}}$ are those the AC is off and locked. These sets get updated, in real time, whenever the AC changes on/off state and are not determined a priori,

$$m[k+1] = \begin{cases} 1, & \mathcal{K}_{\text{lock}}^{\text{on}} : \big[k+1, .., k + t_{\text{locked}}^{\text{on}}\big] \\ 0, & \mathcal{K}_{\text{lock}}^{\text{off}} : \big[k+1, .., k + t_{\text{locked}}^{\text{off}}\big] \end{cases}$$

where $t_{\text{locked}}^{\text{on}}$ is the AC turn-on lock-out duration and $t_{\text{locked}}^{\text{off}}$ is the AC turn-off lock-out duration. The ACs make turn-off requests R_n^{off} between the minimum epoch duration $t_{\min}^{\text{on}}$ and the maximum epoch duration $t_{\max}^{\text{on}}$ according to the exponential distribution function,

$$R_n^{\text{off}}\big(T_{\text{a},n}[k], t_n[k]\big) := \\ \big(1 - e^{-\mu^{\text{off}}(T_{\text{a},n}[k])\Delta t}\big)\big(1 - e^{-\gamma(t_n[k])\Delta t}\big) \tag{3.6}$$

where t_n is the elapsed epoch duration which shall lie between the minimum and maximum epoch durations, and μ^{off} is the temperature-based turn-off rate function defined as,

$$\mu^{\text{off}}(T_{\text{a},n}[k]) = \begin{cases} 0, & \text{if } T_{\text{a},n}[k] \geq T_n^{\max} \\ g(T_{\text{a},n}), & \text{if } T_{\text{a},n}[k] \in (T_n^{\min}, T_n^{\max}) \\ \infty, & \text{if } T_{\text{a},n}[k] \leq T_n^{\min} \end{cases} \tag{3.7}$$

and

$$g(T_{\text{a},n}[k]) = \left(\frac{T_n^{\max} - T_{\text{a},n}[k]}{T_{\text{a},n}[k] - T_n^{\min}}\right) M_{\text{off}} \tag{3.8}$$

where M_{off} is the mean turn-off request frequency, similarly defined as M_{on}, while $\gamma(t_n[k])$ is the time-based turn-off rate function defined by

$$\gamma(t_n[k]) = \begin{cases} 0, & \text{if } t_n[k] \leq t_{\min}^{\text{on}} \\ \left(\frac{t_n[k] - t_{\min}^{\text{on}}}{t_{\max}^{\text{on}} - t_n[k]}\right) M_{\text{off}}, & \text{if } t_n[k] \in (t_{\min}^{\text{on}}, t_{\max}^{\text{on}}) \\ \infty, & \text{if } t_n[k] \geq t_{\max}^{\text{on}} \end{cases} \tag{3.9}$$

The dead-band and lock-out dynamics constrain the turn-off request function as follows,

$$P\big(R_n^{\text{off}} | T_{\text{a},n}[k] \leq T_n^{\min} \wedge t_n[k] \geq t_{\max}^{\text{on}} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{on}} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{off}}\big) = 1, \tag{3.10}$$

$$P\big(R_n^{\text{off}} | T_{\text{a},n}[k] \geq T_n^{\max} \vee t_n[k] \leq t_{\min}^{\text{on}} \vee k \in \mathcal{K}_{\text{lock}}^{\text{on}} \vee k \in \mathcal{K}_{\text{lock}}^{\text{off}}\big) = 0. \tag{3.11}$$

where $\mathcal{K}_{\text{lock}}^{\text{on}}$ and $\mathcal{K}_{\text{lock}}^{\text{off}}$ retain their prior definitions above.

There are several shortcomings to this extended PEM approach. The minimum and maximum epoch parameters, $t_{\min}^{\text{on}}$ and $t_{\max}^{\text{on}}$, for each device are heuristically selected, considering each device's manufacturer-defined lock-out duration and the

thermal dynamics and parameters of the space being cooled by each AC. The lock-out duration could vary across manufacturers while the thermal dynamics and parameters certainly varies across ACs. As previously stated, formally determining optimal values for these parameters is cumbersome and inefficient due to the large parameter space [87]. Hence, an alternative approach eliminating the heuristic parameter selections, while providing similar performance guarantees, is most welcome.

Also, the epoch parameters impact the request probability function and limit availability and flexibility of the ensemble. This is shown in Sections 3.3.3 and 3.5. Furthermore, as we shall see in more details in Section 3.4, the epoch parameters have a non-trivial impact in aggregate model formulations, resulting in increased model complexity and curse of dimensionality. These shortcomings described above are addressed in the new controller design described in Section 3.3.3 below.

3.3.3 New Device-Driven Load Control Strategy

In the new controller, the request-response approach with turn-on and turn-off capabilities is retained while the energy packet concept is removed and request approvals will not be time-constrained. This means the flexible shifting of when and how devices consume energy will be dictated by operational objectives and system dynamics, not arbitrarily defined time-based/packetized concepts and heuristically chosen epochs. Hence, for the new controller, the turn-on request probability functions previously defined in (3.1) - (3.5) are retained, however, the turn-off request probability function is now defined by:

$$R_n^{\text{off}}\big(T_{\text{a},n}[k]\big) := \big(1 - e^{-\mu^{\text{off}}(T_{\text{a},n}[k])\Delta t}\big), \tag{3.12}$$

where μ^{off} maintains the definition in (3.7). Incorporating dead-band and lock-out dynamics constrains the turn-off request function as follows,

$$P\big(R_n^{\text{off}} | T_{\text{a},n}[k] \leq T_n^{\min} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{on}} \wedge k \notin \mathcal{K}_{\text{lock}}^{\text{off}}\big) = 1, \tag{3.13}$$

$$P\big(R_n^{\text{off}} | T_{\text{a},n}[k] \geq T_n^{\max} \vee k \in \mathcal{K}_{\text{lock}}^{\text{on}} \vee k \in \mathcal{K}_{\text{lock}}^{\text{off}}\big) = 0. \tag{3.14}$$

ensuring turn-off requests can only be made at time steps k outside the set of locked time steps $\mathcal{K}^{\text{off}}_{\text{lock}}$ and $\mathcal{K}^{\text{on}}_{\text{lock}}$. This yields a relaxed version of (3.10) and (3.11) such that the turn-off request function is now based only on temperature, with the elapsed epoch parameters and functions of the extended PEM in (3.6), (3.9) - (3.11) no longer needed.

Discarding the time-based epoch and packetized concepts results in a new operational sequence. At each control time step, the new controller shall involve both device level and aggregator level control, which are detailed next:

3.3.3.1 Device Level Operations

Firstly, devices evaluate turn-on/turn-off request probabilities according to the cumulative exponential probability functions defined in (3.1) - (3.5), (3.7), and (3.12) - (3.13). Turn-on request probabilities are computed only when TCLs are off and unconstrained by compressor lock-out, and vice-versa for turn-off request probabilities.

Next, after computing request probabilities, R^{on}_n or R^{off}_n, based on local state information, a bernoulli draw is done to determine whether the device makes a request to the aggregator or not. If the draw is successful, the device notifies the aggregator of its availability and willingness to turn-on/turn-off and if the draw is unsuccessful, no action is carried out till the next time step. Note that the probability of success and rate of requests are primarily determined by the AC temperature state and chosen mean turn-on-request frequency/mean turn-off-request frequency of (3.3) and (3.8). Also, requests can be encoded for privacy and fairness concerns, so no state information or other device identifying information is made known to the aggregator, the aggregator shall simply receive an anonymous turn-on/turn-off request.

Subsequently, if the aggregator approves the turn-on request above, the device shall turn-on. After a turn-on, the device shall then start making turn-off requests (when unlocked and available for compressor-based TCLs) and shall only turn-off either when its turn-off request is approved by the aggregator or when at the upper dead-band limit, $T^{\max}$. This logic also conversely applies to approved turn-off

requests. The on/off duration is no longer heuristically time-constrained as in the extended PEM. If the aggregator denies the device request, no action is carried out by the device till the next control time step.

3.3.3.2 Aggregator Level Operations

The aggregator receives anonymous and asynchronous (because each request is based on local state measurements) device requests. Aggregator approves or denies the request according to the system operator command and overall objective. Knowing the average device power consumption, request approval/denial is simply based on whether the aggregator seeks to increase/decrease the aggregate power consumption of the ensemble. Note, the aggregator can estimate its ramp-up (aggregate power increase) and ramp-down (aggregate power decrease) capacities over a time horizon based on the average device power consumption and the number of turn-on and turn-off requests received over that time horizon.

Next, beyond controller operation sequence, we examine factors underpinning the controller turn-off request redesign in (3.12) - (3.14). There is sufficient correlation between the temperature and elapsed epoch state variables, such that it suffices to use one and not both when formulating request probabilities. For a randomly selected AC (applicable to all ACs) simulated under the extended PEM, the left plot of Fig. 3.2 shows a scatter plot created from the normalized indoor air temperature measurements and the elapsed epoch values during a single energy packet duration that illustrates the strong correlation (-0.979 Pearson coefficient) between temperature and elapsed epoch. Also, for the same AC and same energy packet duration, the right plot of Fig. 3.2 is a logarithmic scatter plot indicating a correlation (0.9678 Pearson coefficient) between the temperature-based turn-off rate function, $\mu^{\text{off}}(T_{\text{a},n}[k])$ (determined by (3.7)) and the time-based turn-off rate function, $\gamma(t_n[k])$ (determined by (3.9)). Following the temperature and epoch correlation described, it can be deduced the temperature-based turn-off request probability function, defined

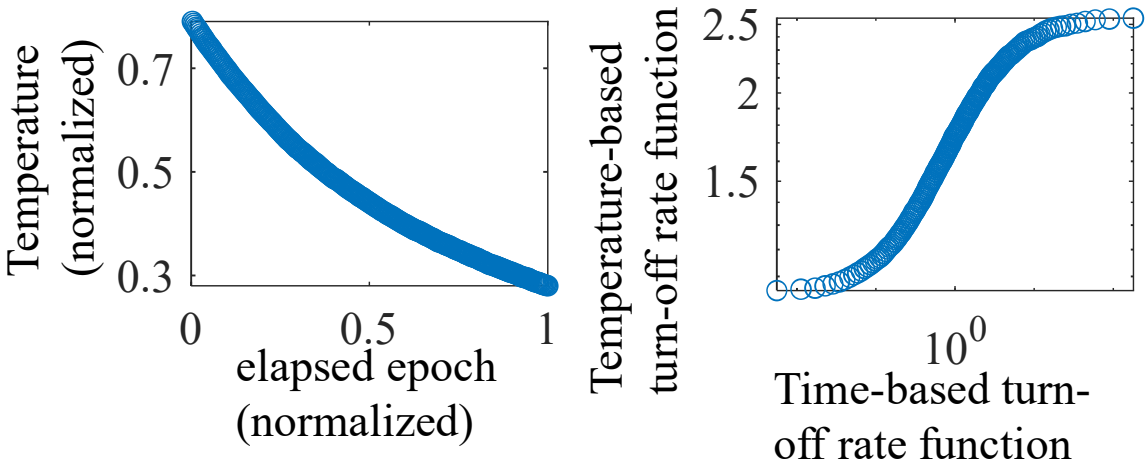


Figure 3.2: Left Plot: Correlation between temperature and elapsed epoch. Right Plot: Log Plot showing correlation between temperature-based turn-off rate function $\mu^{\text{off}}(T_{\text{a},n}[k])$ and time-based turn-off rate function $\gamma(t_n[k])$.

in (3.12), and the time-based turn-off request function,

$$R_n^{\text{off}}\big(t_n[k]\big) := \big(1 - e^{-\gamma(t_n[k])\Delta t}\big), \tag{3.15}$$

are strongly correlated. Furthermore, we can infer that the presence of both rate functions in (3.6) results in a more conservative request rate, illustrated by Fig. 3.3. This is because (3.6) is the product of the temperature-based turn-off probability in (3.12) and the time-based turn-off probability in (3.15), with the latter evolving slower, i.e., only able to attain non-zero values (0^+) later in time, as seen in Fig. 3.4, because of compressor lock-out restrictions, and is thus mostly smaller than (3.12), as shown in Fig. 3.4. The effect is visible in Fig. 3.3 where for a sample AC, we compare probabilities based on using both rate functions as in (3.6) and using only the temperature-based rate function as in (3.12). Note that the normalized elapsed epoch in Fig. 3.3 does not commence from 0 as turn-off request probability is only computed after the AC becomes unlocked and available. Also, the turn-off request probability and turn-off request frequency are functions of the selected M_{off} in (3.9), the higher M_{off} is, the higher the probability and request frequency, and vice-versa.

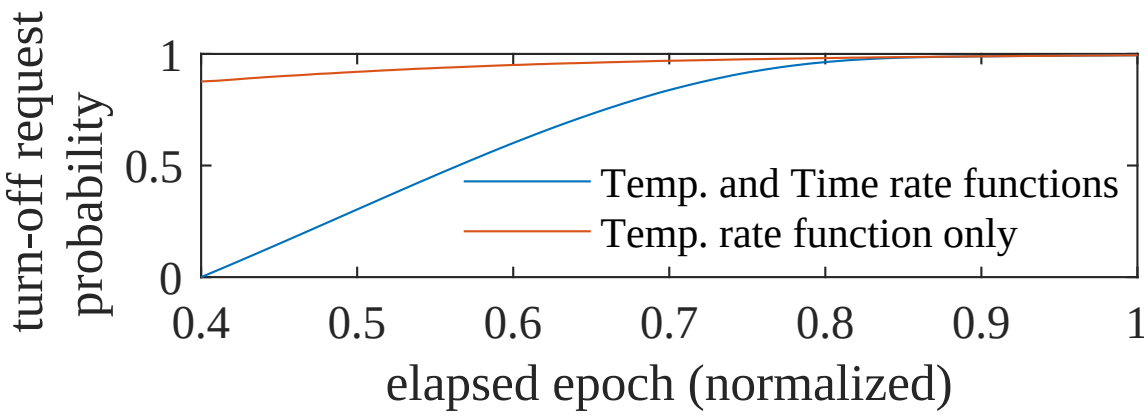


Figure 3.3: Turn-Off Request Probability computation using 2 rate functions ($\mu^{\text{off}}(T_{\text{a},n}[k])$ and $\gamma(t_n[k])$) vs 1 rate function ($\mu^{\text{off}}(T_{\text{a},n}[k])$).

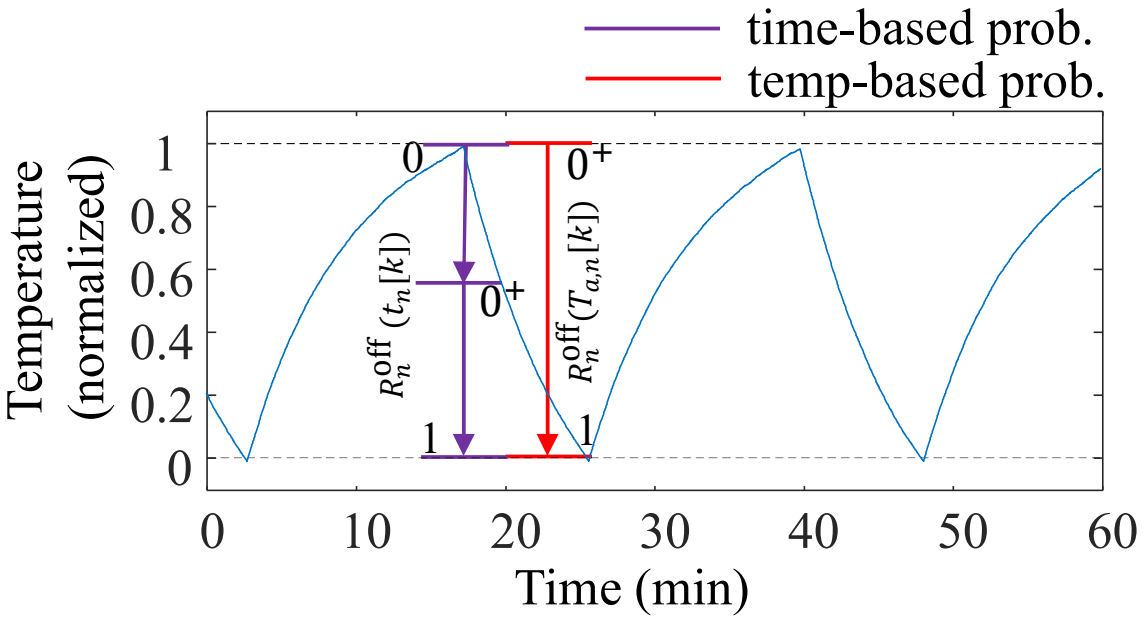


Figure 3.4: Temperature-based and time-based turn-off probability evolution with temperature.

Lastly, analytical feasibility is another motivating factor. The state-space size and complexity blows up with the presence of the epoch variable, hence, the new control design eliminating it to yield a smaller and well defined state-space not dependent on heuristically chosen parameters, is useful. This is discussed in detail in Section 3.4.

To summarize, the following are the fundamental differences between the new controller and extended PEM:

1. Eq. (3.6) is replaced with (3.12) which uses only the temperature-based turn-off probability, we strip off the time-based turn-off probability. Also, (3.14) and (3.13) replace (3.10) and (3.11) for the same reason.

2. The epoch concept is removed. The turn-on duration shall no longer be arbitrarily time-constrained. Thus, the aggregator has increased active control on the downward ramp rate as the passive downward ramp from expiring epochs is now eliminated.

3. 1 and 2 above eliminate all time-based computations in the controller. Hence, there is no longer a need for internal clocks (keeping track of elapsed epoch). Thus, practical implementation is easier and cheaper.

4. Eliminating the epoch variable reduces the state space size and complexity, easing modeling and analysis.

5. The relaxed turn-off request function yields increased flexibility for the new controller, unlocking improved signal tracking ability in the new controller. Section 3.5 provides extensive quantitative and qualitative performance comparisons of extended PEM and the new controller.

3.4 Aggregate Modeling

Ref. [63, 112, 114] developed aggregate model formulations for TCLs under the basic PEM scheme. Ref. [63] developed a model to represent the aggregate behavior of resistive water heaters under basic PEM, [112] developed and validated state bin

transition aggregate models for TCLs and energy storage systems (ESS) and [114] extended the prior model to cater to a fleet consisting of electric vehicles, TCLs, and ESS, while also providing end-user quality of service guarantees. These aggregate models capture the collective behavior of diverse energy resources under PEM but fail to capture the lock-out dynamics of the ubiquitous compressor-based class of TCLs, which [30, 117] showed to be non-trivial.

In this section, we develop an aggregate dynamical system model that approximates the behavior of an ensemble of compressor-based TCLs under the new device-driven control. The compressor lock-out dynamics and turn-off flexibility characteristics of the new controller, not previously captured in past PEM aggregate models is now incorporated in the proposed model. Having a model that provides analytical insights into observed dynamics and lends itself to established estimation, control and analysis techniques in the literature strengthens the potential for adopting this technique/resource in practice. It is also useful for evaluating anticipated performance and expected dynamics for a broad range of scenarios in which individual simulations may be impractical, as well as ascertaining system steady state without simulation.

The aggregate model takes the form of the established system transition model in the literature [31, 130]. This model however differs from those because those transition models were based on aggregator-driven methods, however, the new controller approach is based on a device-driven architecture. The state vector under the aggregator-driven methods represents distributions of the TCL populations receiving probabilistic control commands while in the device-driven approach, TCLs make probabilistic requests but receive deterministic control commands. Hence, the state transition matrix in this work shall incorporate both the turn-on and turn-off request functions of the device-driven controller, as well as the deterministic aggregator control commands.

The model shall approximate the state (temperature state, lock-out state and on/off state) progression and aggregate power trajectory of the TCL ensemble under the new controller. The model is subsequently validated using select case studies, shown in Section 3.5.

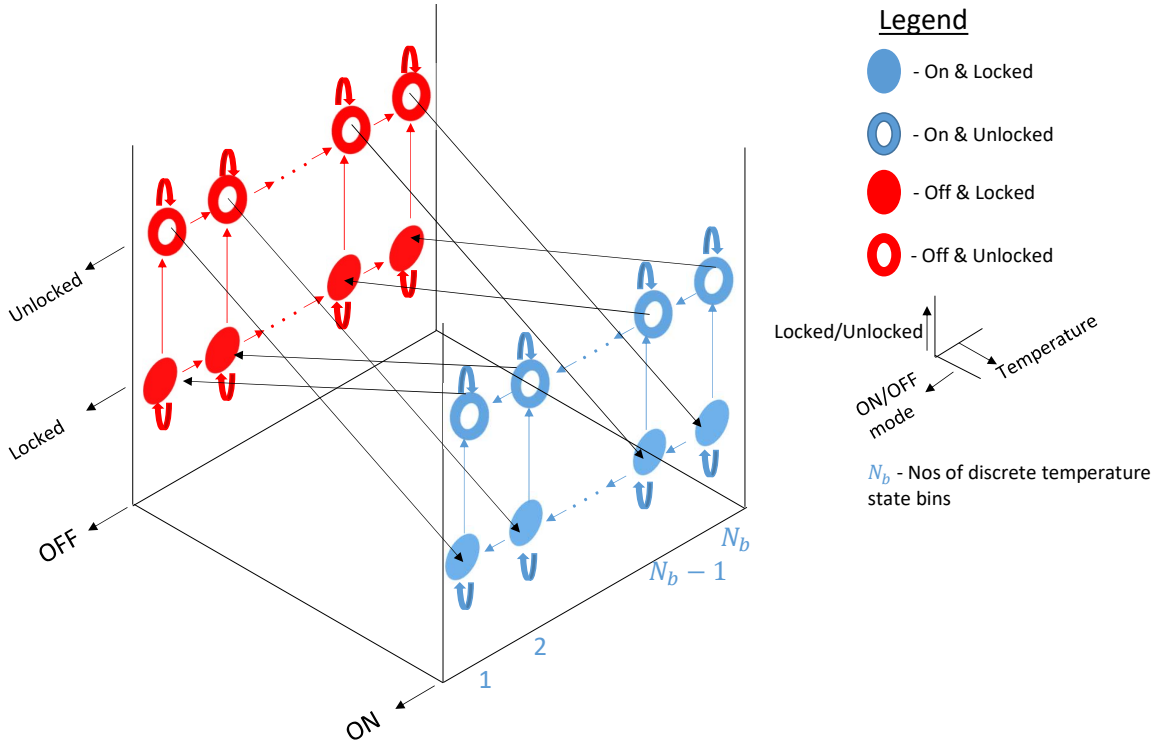


Figure 3.5: State bin transition model for ACs under the new device-driven controller.

We discretize the temperature, lock-out and on/off state space of all TCLs into the same set of "bins", as shown in Fig. 3.5, and the progression of the TCLs across these bins is modeled with transition probabilities. In Fig. 3.5, each circle signifies a bin of TCLs that are either locked/unlocked, on/off and their temperature state progression. The blue solid bins signify the TCLs that are on and locked, the blue hollow bins indicate TCLs that are on and unlocked. Red solid bins are off and locked while red hollow bins are off and unlocked. The temperature state has been discretized into N_b intervals. Horizontal arrows in the figure indicate temperature transitions according to the autonomous thermal dynamics. Vertical arrows indicate a transition from the locked state to the unlocked state. Diagonal arrows indicate transitions between on/off states according to the default thermostat control and/or the new controller. The self-loop arrows on each bin indicate the fraction of TCLs remaining in that bin across the discrete time progression.

Let us define a linear time-varying system:

$$\begin{aligned} x[k+1] &= A(u)[k]x[k] \\ y[k] &= Cx[k] \end{aligned} \tag{3.16}$$

where $x[k]$ is the state vector at time step k, A is the state transition matrix, and a function of control input u, y is the system output and C is the output matrix. The state variable x defines the fraction of the TCL population across the bins in Fig. 3.5. Let us categorize x as follows: $x = [x_1 ... \ x_{N_\mathrm{b}}, x_{N_\mathrm{b}+1} ... \ x_{2N_\mathrm{b}}, x_{2N_\mathrm{b}+1} ... \ x_{3N_\mathrm{b}}, x_{3N_\mathrm{b}+1} ... \ x_{4N_\mathrm{b}}]^T$ where the first N_b bins : $x_1 ... \ x_{N_\mathrm{b}}$ represent the TCLs that are off and unlocked, the next N_b bins : $x_{N_\mathrm{b}+1} ... \ x_{2N_\mathrm{b}}$ represent the TCLs that are on and unlocked, the next N_b bins : $x_{2N_\mathrm{b}+1} ... \ x_{3N_\mathrm{b}}$ represents TCLs that are off and locked and the last N_b bins : $x_{3N_\mathrm{b}+1} ... \ x_{4N_\mathrm{b}}$ represent TCLs that are on and locked.

The state transition matrix, $A(u)$, is formulated as the product of two transition matrices: $A(u) = A_{c,u}A_d$. The matrix A_d models the state transition probabilities according to the autonomous thermal and lock-out dynamics, and default thermostat control dynamics of the TCLs. The matrix A_d, which is the transpose of the Markov Transition Matrix, can either be computed analytically as shown in [31] or it can

be identified for specific outdoor temperatures by counting the bin transitions in the absence of aggregator control. In this work, we assume a constant outdoor temperature for simplicity and because the case studies are over time horizons of one hour, within which a constant outdoor temperature is a reasonable assumption, making A_d time-invariant.

Till this point, the modeling is identical to that in [130], however, the following model definitions which is employed in this work, differ from that of [130] as the TCLs switching dynamics and aggregator control actions here are driven by the request-approval concept.

Matrix $A_{c,u}$ models the state transition probabilities based on the new controller's algorithm and exponential request functions in (3.1) - (3.5), (3.7) and (3.12) - (3.13), and the time-varying control input u defining the fraction of turn-on requests, u^{on} and the fraction of turn-off requests, u^{off}, approved by the aggregator. Hence, $A_{c,u}$ takes the following form: $A_{c,u} = \begin{bmatrix} \mathbf{I} - u^{\text{on}}.diag\big(R_n^{\text{on}}\big) & 0 & 0 & 0 \\ 0 & \mathbf{I} - u^{\text{off}}.diag\big(R_n^{\text{off}}\big) & 0 & 0 \\ 0 & u^{\text{off}}.adiag\big(R_n^{\text{off}}\big) & \mathbf{I} & 0 \\ u^{\text{on}}.adiag\big(R_n^{\text{on}}\big) & 0 & 0 & \mathbf{I} \end{bmatrix}$

where $\mathbf{I}$ is the identity matrix, $diag\big(R_n^{\text{on}}\big)$ and $diag\big(R_n^{\text{off}}\big)$ map the approved turn-on and turn-off requests computed with (3.1) and (3.12) respectively, to the diagonal matrix; $diag\big(R_n^{\text{on}}\big)$ is given by $diag\big((P(\mathrm{R}_1^{\text{on}}), ..., P(\mathrm{R}_{\mathrm{N_b}}^{\text{on}}))\big)$ representing the turn-on request probabilities across bins; $diag\big(R_n^{\text{off}}\big)$ represents $diag\big((P(\mathrm{R}_{\mathrm{N_b}+1}^{\text{off}}), ..., P(\mathrm{R}_{2\mathrm{N_b}}^{\text{off}}))\big)$, the turn-off request probabilities across bins; adiag(.) maps requests in the upper right corner -to- lower left corner anti-diagonal matrix with $adiag(\big(R_n^{\text{on}}\big)$ representing $adiag(\big(\big(P(\mathrm{R}_{\mathrm{N_b}}^{\text{on}}), ..., P(\mathrm{R}_1^{\text{on}})\big)\big)$, and $adiag(\big(R_n^{\text{off}}\big)$ representing $adiag(\big(\big(P(\mathrm{R}_{2\mathrm{N_b}}^{\text{off}}), ..., P(\mathrm{R}_{\mathrm{N_b}+1}^{\text{off}})\big)\big)$. The diagonal sub-matrices model transitions out of the unlocked bins while the anti-diagonal sub-matrices model transitions into the locked bins. We assume for a given time-step, that the autonomous TCL dynamics occur before the controller dynamics. Therefore, the entries in the state transition matrix $A(u)[k]$ now define the overall probability that a TCL under the new controller will transi-

tion from bin to bin, with the form

$$A(u)[k] = \begin{bmatrix} p_{1,1} & \cdots & p_{1,4N_b} \\ \vdots & \ddots & \vdots \\ p_{4N_b,1} & \cdots & p_{4N_b,4N_b} \end{bmatrix}$$

where $p_{i,j}$ is the probability of a TCL transitioning from bin j to bin i in the next time step.

The output matrix has the form:
$C = [0 \ ... \ 0, P^{\text{on}}N_{\text{TCL}} \ ... \ P^{\text{on}}N_{\text{TCL}}, 0 \ ... \ 0, P^{\text{on}}N_{\text{TCL}} \ ...P^{\text{on}}N_{\text{TCL}}]$ where P^{on} is the average power consumption of TCLs in the ON state and N_{TCL} is the total number of TCLs in the ensemble. It is structured in accordance to the state vector x, so that the TCLs in bins $1 : N_{\text{b}}$ and $2N_{\text{b}+1} : 3N_{\text{b}}$ have zero power draws since they represent TCLs that are OFF.

The above model developed for the new controller can be extended for extended PEM as it captures most of the TCL control dynamics in the extended PEM scheme, but it does not capture the epoch concept which is fundamental to extended PEM. The spontaneously occurring state transitions due to expiring epochs/energy packets are not represented in that state transition model as it does not explicitly model packet durations. In order to capture the epoch and packet expiration dynamics for the TCLs, the time state should be discretized (time interval Δt) into a finite set of bins and embedded into the model, resulting in the discrete time and state transition model shown in Fig. 3.6. Similar model definitions employed above can also be used in defining this expanded state space. It can be seen that the number of states grows significantly in this model compared to that of the new controller. While the new controller has a state size of $4N_{\text{b}}$, the typical extended PEM scheme with an epoch of 10 minutes, maximum turn-off duration of 30 minutes and a model discretization time interval of 30 s, results in a state size of $80N_{\text{b}}$.

Hence, a model of the new controller with $N_{\text{b}} = 20$ has a state vector size of 80 while the extended PEM model has 1600. It can be deduced that model identification and computations will be more involved using this extended PEM model, compared

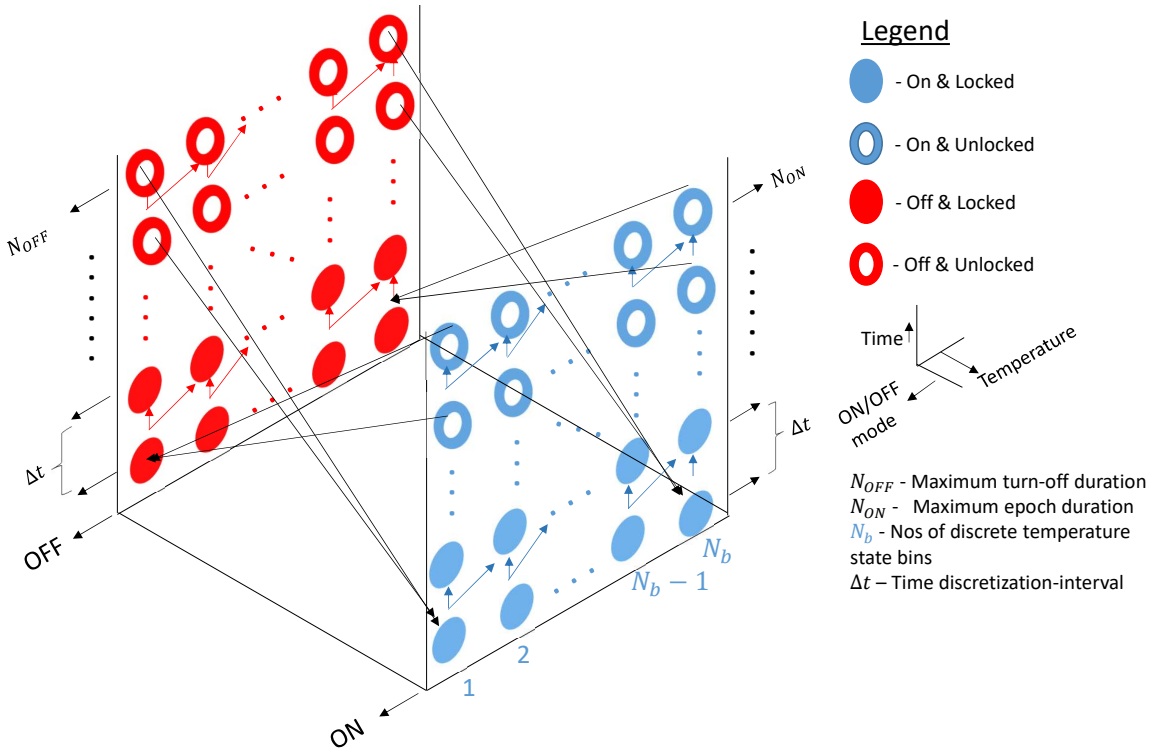


Figure 3.6: State bin transition model for ACs under the extended PEM scheme.

to the new controller model, given the larger state space. Since the new controller offers improved flexibility and better performance, model discussions and validation in this chapter shall focus only on the new controller model.

3.5 Case Study

3.5.1 Simulation Scenarios and Parameters

This section describes the test cases used in evaluating both extended PEM and the new device-driven controller. An objective of the case study is to highlight performance trade-offs or improvements arising from the controller redesign in Section 3.3.3. Five conditions, shown in Table 3.2, were varied between "nominal" and "extreme", to characterize the performance limits of the controllers, as well as issues that may arise from AC coordination under extreme conditions.

Nominal conditions are defined by reference signal amplitudes which are 10% of the ensemble's average power consumption when not controlled (i.e. the baseline power), outdoor temperature of $32.2°C$, a heterogeneous ensemble with AC thermal parameters generated using GridLAB-D [118], with $\pm 20\%$ random variation around each parameter, and perfect communication network characteristics (i.e. no communication delays and no packet losses). Extreme conditions use reference signal amplitudes that are 30% of the baseline power, outdoor temperature of $100°F$, an homogeneous ensemble of AC's with half the thermal capacity used in the nominal condition and generated using GridLAB-D [118], with no variation in parameters, and imperfect communication network characteristics (i.e. delays normally distributed with mean 18 s and standard deviation 3 s, and packet loss across all input commands uniformly distributed between 5 and 10%).

The case studies consider the response of over a thousand (1079) ACs that are controlled to track three reference signals across both nominal and extreme conditions, yielding six cases. Firstly, a sinusoidal signal (1 hr period), secondly, a square signal (inducing a step-like response) and thirdly, a historical PJM RegD signal [121]. The first two signals are slow time-varying signals chosen to better understand the

Table 3.2: Simulation Conditions

Condition	Nominal	Extreme
Reference Signal Amplitude	10%	30%
Outdoor Temperature	$32.2°C$	$37.7°C$
AC Heterogeneity	Heterogeneous	Homogeneous
AC Thermal Parameters	Default [118]	50% of default [118]
Communication Network	Perfect	Imperfect

Table 3.3: Simulation Parameters

Parameters	Value	Parameters	Value
Δt	2 s	Simulation Period	1 h
$t^{\text{on}}_{\min}$	4 min	$t^{\text{on}}_{\max}$	10 min
T_n^{set}	$[20, 24]°\text{C}$	$T_n^{\max} - T_n^{\min}$	$[1, 2]°\text{C}$
T_{o}	37.78°C	M_{off}	1 Hz
M_{on}	$\frac{1}{300}$ Hz	PEM dead-band	$0.8(T_n^{\max} - T_n^{\min})$

system behavior while PJM RegD is a fast time-varying signal offering real-world frequency regulation attributes.

Each AC consumes 2.5 kW on average during regular operation. This study considers houses with single central AC units, but is directly applicable and extendable to houses with multiple AC units given that every AC unit will make its own request. Table 3.3 lists the simulation parameters that were used. As explained in Section 3.3.2, we constrict the temperature dead-band used by the PEM controller, creating a PEM dead-band of width $0.8(T_n^{\max} - T_n^{\min})$ for all ACs, where 0.8 is a design choice that has been found to sufficiently mitigate violations of the AC thermostat dead-bands. Hence, ($T_n^{\max}$, and $T_n^{\min}$) used in (3.4) - (3.14) are 80% of the original AC ($T_n^{\max}$, and $T_n^{\min}$) under default thermostat control. Also, in the extreme cases, aggregator request approvals rates are scaled by a factor (akin to control gain) of 0.2, to mitigate oscillations arising from the communication delays of extreme conditions.

Table 3.4: Controller performance results across nominal cases

	Case 1A	Case 1B	Case 2A	Case 2B	Case 3A	Case 3B
Controller Type	Old	New	Old	New	Old	New
Signal Type	Sinusoidal	Sinusoidal	Step	Step	Frequency Regulation	Frequency Regulation
Regulation Capacity (kW)	N/A	N/A	N/A	N/A	107	107
NRMSE (%)	0.54	0.38	0.72	0.62	0.62	0.52
PJM Accuracy Score (0-1)	N/A	N/A	N/A	N/A	0.99	0.99
PJM Delay Score (0-1)	N/A	N/A	N/A	N/A	0.97	0.97
PJM Precision Score (0-1)	N/A	N/A	NA	N/A	0.99	0.99
PJM Composite Score (0-1)	N/A	N/A	N/A	N/A	0.98	0.98
Mean Availability (%)	58.93	68.34	59.12	67.34	57.75	64.73
Mean Ramp Down (%)	12.69	23.69	12.84	23.42	11.89	22.03
Mean Ramp Up (%)	1.36	2.26	1.22	1.92	1.23	1.77

Table 3.5: Controller performance results across extreme cases

	Case 4A	Case 4B	Case 5A	Case 5B	Case 6A	Case 6B
Controller Type	Old	New	Old	New	Old	New
Signal Type	Sinusoidal	Sinusoidal	Step	Step	Frequency Regulation	Frequency Regulation
Regulation Capacity (kW)	N/A	N/A	N/A	N/A	561	561
NRMSE (%)	13.19	2.85	10.89	8.27	26.59	23.27
PJM Accuracy Score (0-1)	N/A	N/A	N/A	N/A	0.69	0.76
PJM Delay Score (0-1)	N/A	N/A	N/A	N/A	0.82	0.83
PJM Precision Score (0-1)	N/A	N/A	N/A	N/A	0.81	0.83
PJM Composite Score (0-1)	N/A	N/A	N/A	N/A	0.77	0.81
Mean Availability (%)	33.04	48.83	35.92	46.97	34.93	39.23
Mean Ramp Down (%)	16.55	37.26	18.70	35.03	14.85	27.07
Mean Ramp Up (%)	0.84	1.18	0.98	1.26	1.36	1.35

3.5.2 Simulation Results

Tables 3.4 and 3.5 summarize the performance across all six cases for both the new controller and extended PEM. Cases 1, 2 and 3 represent nominal conditions while 4, 5 and 6 represent extreme conditions. Cases 1A through 6A represent extended PEM controller evaluations while 1B through 6B are the new controller evaluations. Tables 3.4 and 3.5 give the normalized root mean square error (NRMSE) between the aggregate power and the power reference, normalized by the baseline power; an industry performance metric called PJM performance scores [122, 131] are also provided; the mean availability, which is calculated as the average percentage of available devices over the simulation period is also provided. An AC is considered available when it is unlocked (and within the PEM temperature dead-band for the extended PEM controller), and hence available to participate and make turn-on or turn-off requests. The mean ramp up/down flexibility which is calculated as the average of the percentage of devices making turn-on and turn-off requests over the simulation horizon is also provided in the tables. These availability and mean ramp up/down parameters indicate the flexibility the aggregator has in increasing/decreasing aggregate power by approving turn-on/turn-off requests.

Comparing the results across all six cases for both controllers, we see the new controller has lower NRMSEs when compared against the extended PEM; In the nominal cases, we see an ≈30% factor decrease in the NRMSE of Case 1A from 0.54% compared to 0.38% for case 1B; similarly under extreme cases, we see NRMSEs decrease up to 78.39% when comparing Case 4A and Case 4B. Table 3.4 and 3.5 show there is improved signal tracking performance in all cases from the new controller. We also see similar trends with the availability and ramp up/down numbers, affirming increased flexibility with the new controller. The increase in NRMSE under extreme conditions is expected, as the availability and ramp up/down flexibility are limited by the conditions imposed which push the controller limits. Fig. 3.7 shows the reference tracking results for the new controller under extreme conditions. We see from the top two plots of Fig. 3.7 that the performance under the step signal input is noticeably worse than under the sinusoidal input scenario because of the steep step change in

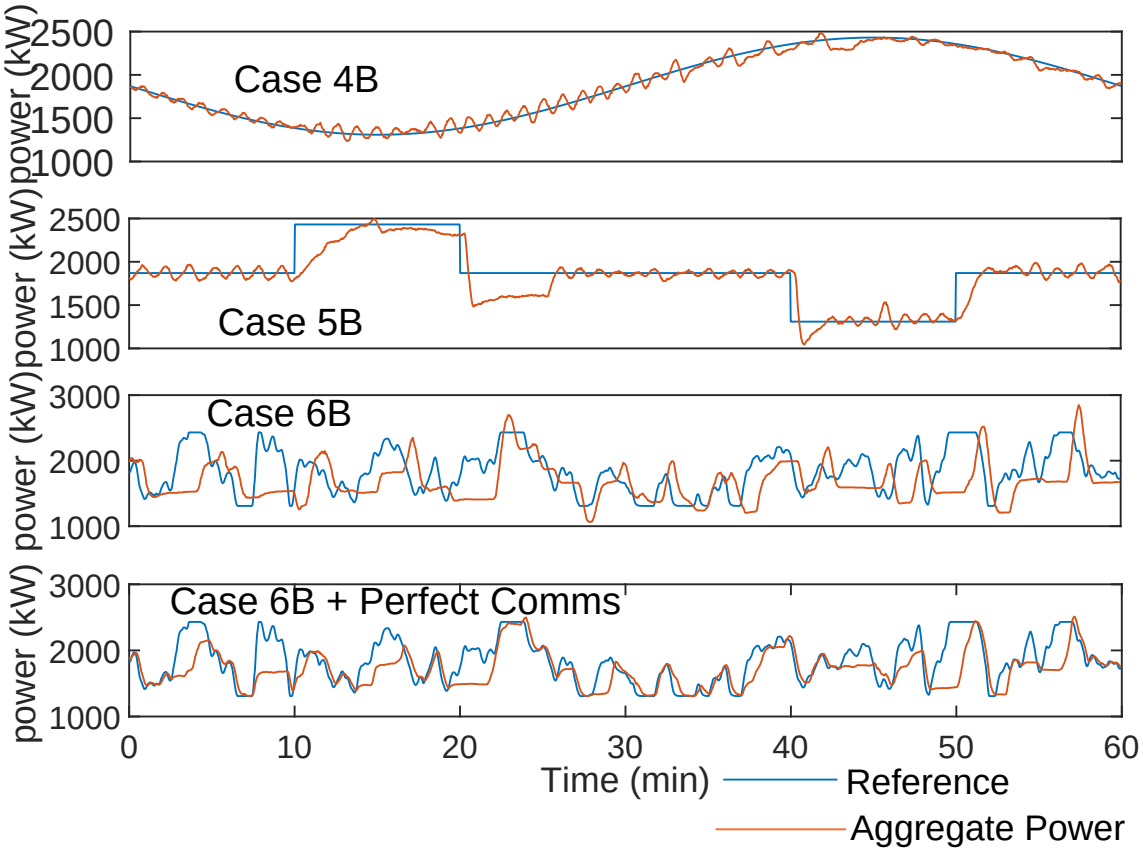


Figure 3.7: Top 3 plots: reference tracking for cases 4B, 5B and 6B respectively; bottom plot: case 6B with perfect communication network.

the signal as opposed to the slowly varying nature of the sinusoidal signal. The ramp up/down flexibility is simply incapable of satisfying the steep step change, an ≈ 600 kW instantaneous step up/down under the extreme step response conditions. This limitation is more pronounced in the faster varying frequency regulation signal of case 6, as the third plot of Fig. 3.7 shows the steeper ramp up/down step-like changes of over 1000kW which pushes the ensemble to its limits. Hence, we see the worst NRMSE values under the frequency regulation signal and the best NRMSE values under the sinusoidal signal input.

PJM requires composite scores greater than 0.75, we see from Tables 3.4 and 3.5, that the controllers qualify for market participation in all cases. However, given the noticeable sharp decline in the performances under extreme conditions compared to the nominal conditions, circumstances such as these should see the aggregator be more conservative with regulation capacity bids, in order to comfortably surpass the qualifying criteria and have sufficient availability and flexibility. It is also possible to ameliorate the effects of the extreme condition scenario by increasing the turn-on request frequency. For example, increasing the turn-on request frequency of Case 6B from 1/300 Hz to 1 Hz while keeping every other parameter constant results in a 1.41% reduction in NRMSE to 21.86%. The tracking performance, availability and flexibility improve with an increase in the turn-on/turn-off request frequencies; however, this increases the communication requirements for the control scheme with requests being made more often. The effects of the communication delays and losses are particularly visible in the third and bottom plots of Fig. 3.7. Keeping all conditions constant but adjusting only the communication network characteristics from extreme to nominal, we obtain the bottom plot of Fig. 3.7 where the signal tracking lag is now gone and performance shortcoming is because of the other extreme condition restrictions. It has a NRMSE of 15.18%, PJM accuracy, delay, precision and composite scores of $0.78, 0.91, 0.91$ and 0.87 respectively and mean availability, ramp down/up percentages of $24.9\%, 19.39\%$ and 2.44% respectively. It is observed that the imperfect communication condition has the most significant effect on performance among all the conditions varied. This demonstrates the need for delay-awareness to be incorporated into the controller design.

3.5.3 Aggregate Model Results

Using the control inputs and state information from the case study simulations run a priori, the aggregate model developed in Section 3.4 is validated against simulation. The model is used to approximate the aggregate power dynamics of the ensemble with the matrix multiplications of (3.16). The model's output is then compared against simulation, which used the equivalent thermal parameter model in Section 3.3.1 and device-driven control of Section 3.3.3.

Let $x_a[k]$ represent the actual state distribution (for time step k), obtained from the simulation, of TCLs across state bins described in Section 3.4 above. Likewise $u_a[k]$ denotes the control actions, i.e. request approvals/denials, taken by the aggregator at time step k. The model estimates the subsequent state progression from time step k according to $x_e[k+1] = A(u_a[k])x_a[k]$ and the aggregate power consumption estimate is $y_e[k+1] = Cx_e[k+1]$ where state matrices A and C retain their definitions from Section 3.4.

The simulation assumes aggregator control actions issued at time step k result in AC turn-on/turn-off at time step $k+1$. Therefore, a one step ahead model prediction, that is, $x_e[k+1]$ is based on past state information $x_a[k]$ and control inputs $u_a[k]$, is used for the validation, shown in the top plot of Fig. 3.8. The model does not perfectly match the simulation, which is expected, as the aggregate model is an approximation of the expected ensemble behavior. The model will prove useful as [132, 133] show improvements in dynamics and performance for PEM when using model predictive control in tracking automatic generation control signals for systems with energy-constrained resources, a future study consideration for the new controller. However, the larger the prediction horizon, the less accurate the model predictions are, as illustrated in Table 3.6, which shows the NRMSE between the simulation and model for different prediction time-steps. Given knowledge of the initialization conditions, $y_e[1] = Cx_a[1]$. For one-step-ahead prediction, $x_e[2] = A(u_a[1])x_a[1]$ and $y_e[2] = Cx_e[2]$. Similarly, $x_e[3] = A(u_a[2])x_a[2]$ and $y_e[3] = Cx_e[3]$.The pattern remains the same subsequently with model outputs $y_e[k]$ being determined by actual state information $x_e[k-1]$.

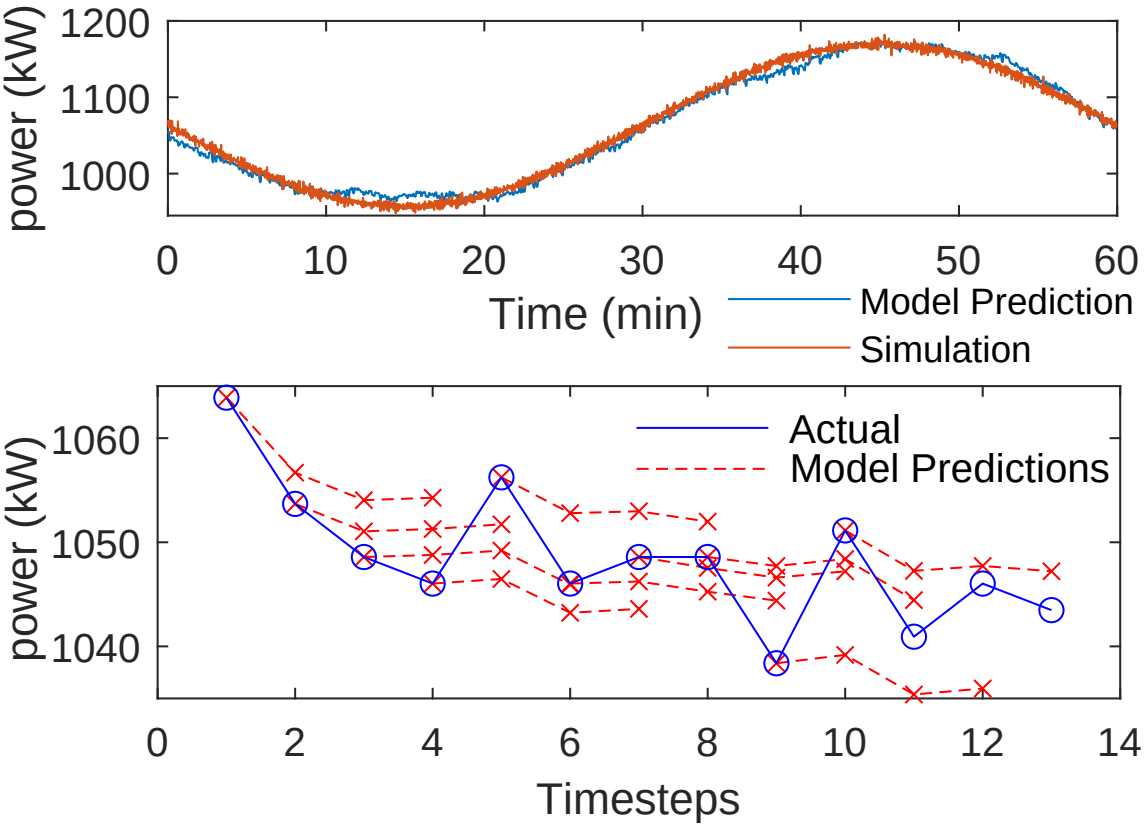


Figure 3.8: Case 1B: simulation vs 3-steps ahead aggregate model prediction.

Table 3.6: Model Prediction Accuracy

Case	1-Step(2 s)	3-Steps (6 s)	10-Steps (20 s)	30-steps(60 s)
1B	0.95%	0.94%	1.12%	2.10
2B	0.96%	1.11%	1.74%	3.43%
3B	0.90%	1.53%	3.75%	7.33%

For three-steps-ahead prediction, $x_e[2]$ retains it definition above but $x_e[3] = A(u_a[2])x_e[2]$, $x_e[4] = A(u_a[3])x_e[3]$ and $y_e[4] = Cx_e[4]$. The model estimate output y_e at time step 4 is based on the actual system state information, $x_a[1]$, obtained three time steps ago. The pattern remains the same subsequently with model outputs $y_e[k]$ being determined by actual state information $x_e[k-3]$. This approach remains the same for any multiple steps ahead prediction.

The bottom plot of Fig. 3.8 shows three-steps ahead predictions for the first ten timesteps, demonstrating how and why prediction accuracy decreases with larger prediction horizons due to the compounded dead reckoning error.

3.6 Summary and Conclusion

In this work, we design a new device-driven controller by eliminating the epoch/ packetized concepts of PEM and utilizing only one rate function in turn-off request formulation, unlocking both practical and analytical benefits. We demonstrate that this leads to improved controller performance, availability and flexibility. The controller was evaluated under nominal and extreme conditions, that highlight the controller limitations against large reference step changes. Also, the impact of ambient conditions on availability and flexibility, and how imperfect communication network degrades controller performance is demonstrated.

Nonetheless, we show its possible for an ensemble of around 1000 ACs to satisfactorily offer frequency regulation even under extreme conditions. Also, an aggregate system model approximating the behavior of an ensemble of ACs under the control schemes was also developed and validated with a one-step ahead model prediction test. Additional uses of the model are also discussed.

CHAPTER 4

A Hardware-in-the-Loop Experimental Testbed using Air Conditioners for Grid Balancing

Driven by the need to offset the variability of renewable generation on the grid, development of load control is a highly active field of research. However, practical use of residential loads for grid balancing remains rare, in part due to the cost of communicating with large numbers of small loads and also the limited experimentation done so far to demonstrate reliable operation. To establish a basis for the safe and reliable use of fleets of compressor loads as distributed energy resources, we constructed an experimental testbed in a laboratory, so that load coordination schemes could be tested at extreme conditions. This experimental testbed was used to tune a simulation testbed to which it was then linked, thereby augmenting the effective size of the fleet. Modeling of the system was done both to demonstrate the experimental testbed's behavior and also to understand how to tune the behavior of each load. Implementing this testbed has enabled rapid turnaround of experiments on various load control algorithms, and year-round testing without the constraints and limitations arising in seasonal field tests with real houses. Experimental results show the practical feasibility of an ensemble of small loads contributing to grid balancing.

4.1 Chapter Introduction

New deployments of renewable energy resources have increased rapidly over the last decade, driven both by new policies to address climate change and by economics as the price of renewable generation rapidly decreases [134]. Given that the fastest growing forms of renewable generation – solar and wind – are variable and intermittent, more fast-response storage and load control capacity will be needed to maintain a reliable grid as the generation mix evolves with a reduced fraction of schedulable/controllable generation. Load coordination, however, is difficult because different types of electrical loads have different characteristics and customer requirements, so that the reliability, consistency, and performance of aggregations of loads is uncertain. For load aggregations to gain acceptance as grid balancing resources and avoid liability with both customers and system operators, it is necessary to test and debug control schemes to failure in an experimental environment before moving to field testing with real houses.

In this chapter we consider aggregations of air conditioners (ACs) that switch on/off to maintain a temperature near a setpoint. Their power consumption can be controlled by sending switching commands or temperature settings directly to the thermostats. One reason for selecting ACs is because of their ubiquity, with ACs present in almost 90% of single-family homes in the U.S. [135]. Another reason is that residential ACs consume more energy than other appliances, so they have a larger impact on the electricity supply-demand balance if their consumption can be shifted in time [136]. ACs and other compressor-based loads are more challenging to control than resistive loads like electric heaters, because compressors have a "lockout" time between when they are shut off and can be turned on again, to guard against short-cycling. Short-cycling risks damaging the compressor or causing it to stall and draw a high current until it overheats or trips a circuit breaker. ACs also display a variable power draw as they start up and reach equilibrium temperatures on the hot and cold heat exchangers. Finally, the AC's compressor draws a transient inrush current (significantly larger than normal operating current) when switched on, and the impact on transformers and distribution circuits of multiple ACs switching on at

exactly the same time should be considered.

We constructed an experimental testbed at Los Alamos National Laboratory to study the effects of load control on residential window ACs in nearly constant ambient conditions. The experimental testbed consists of 20 single-zone model houses of nearly identical construction, situated within a warehouse. To ensure that these models were representative of real houses, initial open-loop experiments were run and validated against data from 47 homes in Austin, TX supplied by Pecan Street, Inc. [137]. Because even 20 ACs represents only a small aggregation, the experimental testbed was designed to optionally interact with a simulation of hundreds to thousands of virtual houses to boost the effective size of the testbed. Several closed-loop feedback control algorithms were tested to demonstrate the aggregation's ability to provide frequency regulation. Confidence in these integrated hardware-in-the-loop (HIL) experiments was established by comparing results from the experimental houses to the virtual houses to check that the control algorithms did not favor either set. Additional features of the testbed include a programmable heat source, to emulate changes in occupancy of the house, and the ability to tune the natural cycling period of ACs in the house with minor physical modifications.

The experimental testbed described here is one of a very small number of laboratory experiments that have been built at this scale, e.g., [138–141]. In contrast to our testbed, [138] focused on refrigerators; specifically, the authors investigate the ability of refrigerators to provide frequency regulation via experiments on 25 refrigerators. Refs. [139–141] focused on commercial rather than residential buildings; [139] explores demand response from a commercial refrigeration system; and [140, 141] used an experimental facility [142] to demonstrate frequency regulation with a commercial HVAC system.

In summary, the contributions of this chapter are fourfold. First, we design, develop, and validate an experimental testbed for testing AC load control strategies. Second, we develop an extended equivalent thermal parameter (ETP) model of ACs, which captures more of the salient characteristics of these devices, e.g., the cooling time lag. Third, we use the testbed to test three load control methods over a wide range of scenarios, using both physical and simulated ACs, demonstrating

the usefulness of the testbed and the ability of ACs to provide frequency regulation. Fourth, we highlight the opportunities and challenges associated with testing AC load control strategies through HIL experiments. These findings may be of interest to researchers and practitioners exploring methods to test/benchmark load control strategies before pushing them to the field. These findings also highlight some of the inherent challenges in using AC aggregations for grid balancing, and some techniques for overcoming those challenges.

In Section 4.2, we describe the HIL architecture, and the design and features of our experimental testbed. Section 4.3 describes HIL testbed validation. The results of our HIL experiments are provided in Section 4.4, and we conclude the chapter in Section 4.5.

4.2 Hardware-in-the-loop Architecture

The HIL experiment architecture is shown in Fig. 4.1. An aggregator running a load control algorithm receives a desired aggregate power command from, e.g., the system operator, and sends control commands to an AC aggregation. Our testbed differentiates between the experimental testbed (20 physical model houses representing the real plant), which receive the commands via a data acquisition (DAQ) and control system, and the simulation testbed (hundreds to thousands of simulated houses representing the virtual plant), which serves to boost the number of ACs in the aggregation. ACs within the simulation testbed are henceforth referred to as "virtual ACs" and ACs within the experimental testbed are henceforth referred to as "experimental ACs." Control commands and measurements are transmitted back to the aggregator via a communication system. The following subsections detail each component shown in Fig. 4.1.

4.2.1 Experimental Testbed Design

Although our model houses are not scaled from any particular house, they exhibit salient features found in typical houses, as we will show in Section 4.3. Also, although some houses have two or more cooling zones, the testbed model houses are single-

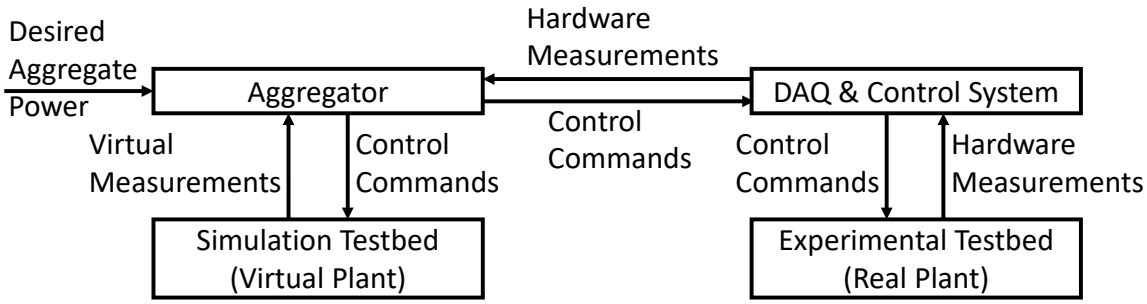


Figure 4.1: Hardware-in-the-Loop experiment architecture.

zone units. Limiting the models to single-zone systems is not especially restrictive, as many single-family houses have single-zone central air conditioning. Also, window ACs are found widely in single rooms of houses, apartments, hotel rooms, and small offices.

The model houses have three subsystems: (1) the experimental AC, (2) the house envelope and environment, and (3) the thermal properties, including the heat source and circulation fan. A sketch depicting the design and a photograph of testbed are given in Fig. 4.2. We next describe each subsystem.

4.2.1.1 Air Conditioners

The ACs used in this testbed are single-speed. Although dual-speed and variable-speed ACs are more efficient and provide improved user comfort, they are not yet dominant in the U.S. [143]. To accommodate as many model houses as possible in the space available, we selected 5000 BTU/h window ACs, which are the smallest capacity commonly sold. The model selected contained a mechanical thermostat, and the thermostat was easily replaced with a relay controlled externally by the DAQ computer. Also, the ACs have a single speed fan, which was turned on all the time both to provide additional air circulation in the box and also to distinguish its power consumption from that of the compressor. The ACs are equipped with a passive element that interrupts power if the compressor stalls and overheats, but adjustable switching delays or lockout periods were imposed in the DAQ software to prevent short-cycling the compressor. Because the built-in mechanical thermostat

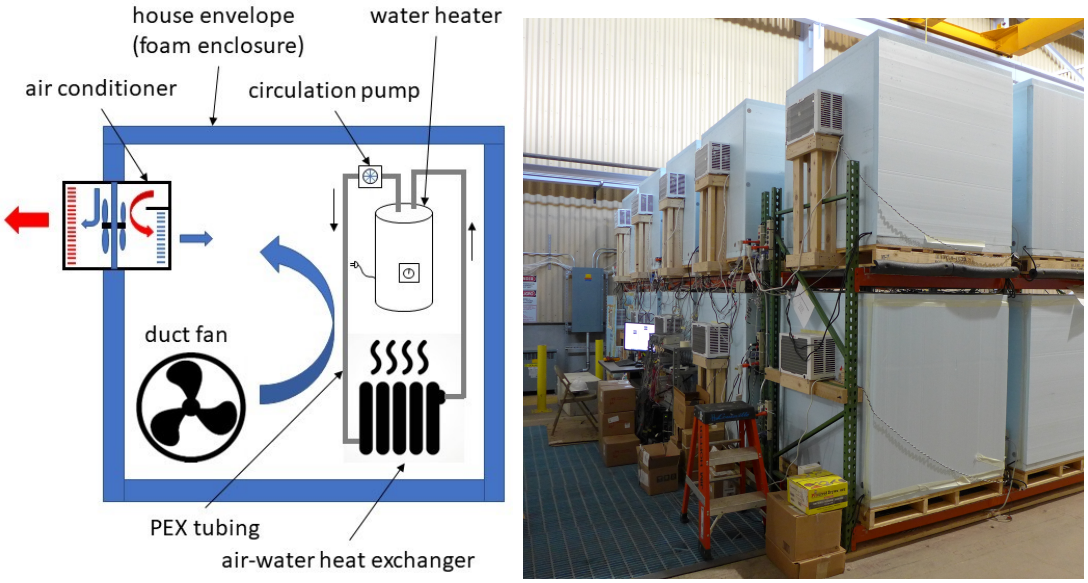


Figure 4.2: Left: Model houses are constructed inside a four foot cubic foam box. The internal heat source consists of a hydronic loop with an electric water heater. A duct fan forces air through the heat exchanger and mixes the room air in addition to the AC's fan. Right: The 20 units are stacked in pallet racks, minimizing cable lengths and overall footprint.

was not used, the DAQ and control system provided the thermostat function in software, switching based on temperatures measured by negative temperature coefficient thermistors in each house.

4.2.1.2 House Envelope and Environment

The insulating boundary of each house was a 1.2 m cube foam enclosure constructed from 5 cm thick extruded polystyrene foam boards with an R-value of 10. The humidity of the air is normally low in Los Alamos, NM with typical values around 10% relative humidity. Relative humidity in the lab was measured with a sensor, but the humidity was not controlled. No adjustments to the data were needed to account for humidity as no condensate was pulled from the air.

4.2.1.3 Thermal Properties

In real houses, solid materials exchange heat with the air and significantly affect the temperature dynamics through their heat capacities. The model houses contain electric point-of-use water heaters, of either 20 or 30 gallon capacity, and this water represents the solid heat capacity normally comprised by the solid walls and furnishings. By adjusting the heat exchange between water and air, one may effectively make the model house appear to have a higher heat capacity and thus cycle the compressor on/off less frequently, given a fixed internal heat load. In this way, the model can behave like a real house despite its smaller size, at least in terms of its AC cycle durations. The air and water are thermally coupled through a 0.3 m square tube-and-fin heat exchanger, and water is circulated from the water heater through cross-linked polyethylene tubing and a pump to the heat exchanger and back to the water tank. This is similar to the hydronic loop of a residential radiant heat system. A single-speed duct fan is placed next to the water-air heat exchanger to flow air through it and mix the air throughout the small box. Were no heat generated in the heater, pump, and fan, the water would simply contribute to the apparent heat capacity in the house.

The water circulation pump and the duct fan are fixed electric loads enclosed in

the house, and they run constantly during experiments. They are therefore constant internal heat loads. In contrast, the water heater is used as a programmable heat source, with the amount of heat injected adjusted by a solid-state power controller. For each on-off cycle, the AC must remove all heat generated in the house during that cycle, independent of the total heat capacity of the air and water.

Each water heater has its own power controller circuit, which allows for each house to have different internal heat gains. The DAQ computer controls all the heating rates and it also allows the operator to add randomness to the heating rate (if needed to disturb phase-locking or synchronization of houses) or to make the heating rate vary over time to emulate occupant usage patterns or diurnal usage variation.

The assembly was constructed in a compact, two-level arrangement shown in Fig. 4.2. Though the construction of each model house is nearly identical, we found that the units tended to have different cooling loads and cycling behavior. This is likely due to a variety of factors, including variation in pumps, fans, and ACs; differences in water heater sizing, differences in envelope construction, differences in the location of the units in the warehouse, component positioning within the units, and differences in the location of the thermistor within the units. The last of these has a large impact.

4.2.2 Simulation Testbed Design

The simulation testbed is described in detail in [56, 98], where we also describe its integration with a field testbed in Austin, TX. Here we briefly describe its main features. The testbed, implemented in MATLAB, incorporates high-fidelity ETP AC models [118] validated against data from homes in Austin, TX [137]. We model AC active power consumption dependence on outdoor temperature, reactive power consumption dependence on voltage, as well as inrush current resulting in high active/reactive power draws that may occur for very short time durations after the compressor turns on. The simulation testbed also enables the analysis of system-wide distribution network effects of load control strategies by assigning each AC to a node of a taxonomy distribution feeder from [144]. GridLAB-D [118] is used as

a power flow solver. Aggregator-to-AC communication networks are modeled with configurable links, package drop rates, and delays.

4.2.3 Aggregator's Load Control Algorithms

We tested three previously-developed load control algorithms. Each has the same objective of regulating the aggregate power consumption of the ACs to match a reference power command; however, each controller differs in architecture and algorithm. All controllers were implemented in MATLAB.

The first controller is a proportional-integral (PI) controller, which uses the system output (aggregate AC power consumption) to compute the reference tracking error, and subsequently the next control input according to standard PI control design [145].

The second control approach is the Markov model-based probabilistic control described in [31]. The benefit of this control scheme is that it uses a model of aggregate AC dynamics to predict AC power consumption one step ahead. It uses broadcast control and can incorporate a state estimator to simplify the design and reduce the cost of the communication network. We use the version of the controller that incorporates delayed temperature dynamics from [146].

The third control approach is the device-driven extended Packetized Energy Management (PEM) strategy described in [30]. PEM, originally introduced in [61], allows devices to make a request to turn on for fixed-duration periods to consume a fixed amount of energy. The load aggregator coordinating the devices approves or denies the requests to track the power reference command. Ref. [30] extended this approach to accommodate compressor-based devices like ACs by including flexible turn-on periods and turn-off requests.

4.2.4 Data Acquisition and Control System

The DAQ and control system serves as a communication intermediary between the aggregator's load control algorithms and the experimental testbed, as shown in Fig. 4.1. The system encodes and sends measurements from the experimental testbed

such as house temperature and AC power consumption to the aggregator, while also processing the aggregator's control commands, both translating the command into AC-specific switching signals and checking whether or not each signal is feasible given the current on/off and lockout states of that AC. Only feasible commands are transmitted.

The DAQ and control system is implemented in National Instruments LabView. The entire array of model houses, the laboratory sensors, and the power meter monitoring of each house are managed by a single computer system. The LabView software includes a simple TCP server that allows it to communicate with other programs on the same computer. Since the AC states for the experimental houses were measured in real time, the DAQ and control system clock, contained in the LabView experiment, determined the latching of time steps, such that the aggregator's controller would step along with the physical experiment.

4.2.5 Communication System

Communication and data exchange of control commands and virtual measurements between the aggregator and simulation testbed were implemented in MATLAB. However, the communication and data exchange between the aggregator and the experimental ACs is achieved via the DAQ and control system using TCP server port connections. Switching control commands from the aggregator are first encoded as digital signals and then further encoded as JavaScript Object Notation (json) strings. This is all done in MATLAB, and sent to a specific TCP port address, i.e., the receiving TCP port address of the DAQ and control system. Likewise, the aggregator receives encoded measurements from the experimental ACs via the DAQ and control system, at its own specific TCP port address. Measurements are subsequently parsed by decoders and data filters built into the aggregator's controller to detect and flag any corrupt data (e.g., infeasible state measurements) before being passed on to the aggregator's control algorithm.

4.3 HIL Testbed Validation

Prior to running controlled HIL experiments, we validated the experimental and simulation testbeds. Specifically, we sought to demonstrate that the experimental testbed characteristics and behavior were consistent with real systems and the simulation testbed sufficiently replicated behaviors and phenomena seen in the experimental testbed.

4.3.1 Experimental Testbed Validation

We conducted open-loop experiments and validated against the same data used to validate the simulation testbed, specifically, one-second interval residential AC submetering data from 47 homes in Austin, TX [137]. We compared the duty cycles, period of cycles, inrush power to steady state power ratio, and variation in power consumption when ACs are on. Additionally, we compared the correlation between power consumption and both outdoor temperature and voltage. For brevity, here we report our comparisons of the variation in power consumption and the power-temperature correlations; full details are available in [147].

The validation experiments were performed across several temperature setpoints, with different temperature deadbands, with different heating rates, and under different ambient conditions. No phase-locking, coupling, or coincidence in the on/off cycling of units was observed despite the relatively close proximity of the units to each other, so they acted as independent oscillators, i.e., there were no groups of two or more units that always cycled on and off at the same time.

For validating the variation in power consumption when ACs are on, we calculate the fractional power variation, defined as the ratio of the difference between minimum and maximum power, and the maximum power. For each on portion of a cycle, we find the minimum and maximum power after discarding the first and last 5 s of data to ensure removal of transients. Histograms of the power variations, for both the experimental ACs and the real ACs, are shown in Fig. 4.3. For the experimental ACs, the mean power variation is $10 \pm 3\%$ whereas for the real ACs, it is $8.7 \pm 9.3\%$. The larger mean variation of the real ACs was skewed by two ACs with high variation

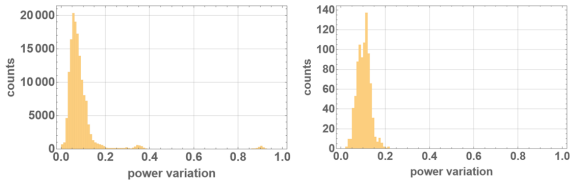


Figure 4.3: Fractional power variation when the ACs are on for the experimental ACs (right) and real ACs (left). The distributions approximately match.

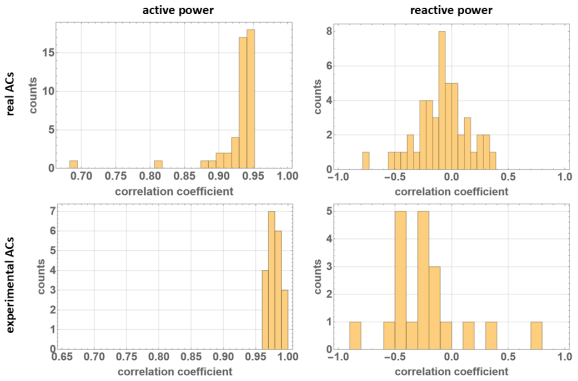


Figure 4.4: Active (left) and reactive (right) power to temperature correlation plots for the real ACs (top) and the experimental ACs (bottom). Note the difference in x-axis scale of the histograms for reactive power.

in some cycles. Nonetheless, the power variation distribution of the experimental ACs approximately agrees with that of the real ACs.

Fig. 4.4 shows histograms of the observed active and reactive power vs. temperature correlations for both the experimental and real ACs. For each AC, we extract power and temperature data from when the AC is on, and compute the scalar correlation coefficient (-1,1) between active power and temperature, and reactive power and temperature. We find that the temperature dependence of power consumption is in fair but not perfect agreement between the experimental and real ACs. This is because the ACs in the testbed are of different models and capacities, and there are no adjustable parameters in the comparison, so perfect agreement could not be expected. The mean fractional change in active power is around $1.36\%/^{\circ}C$ for the experimental ACs and $2.12\%/^{\circ}C$ for the real ACs.

Note that the count difference between the real and experimental data observed in Figs. 4.3 and 4.4 is a function of the limited number of ACs in the experimental dataset (20) versus the real dataset (47) and the shorter time over which the experimental ACs were observed.

4.3.2 Simulation Testbed Validation

Seven open-loop experiments were conducted to verify that any phenomena which are observable in the experimental testbed are also captured and seen in the simulation testbed. The experiments were formulated to induce atypical phenomena, if any, that may not often be observed in practice. For brevity, here we report/discuss comparisons only in terms of power; full details are available in [148].

The first experiment ran the ACs without external control to explore their open-loop power consumption. The second experiment explored synchronization behavior by controlling all ACs to maximize (or minimize) aggregate power consumption until synchronization, and then releasing control to observe the effects of synchronization on aggregate power variation and the time to desynchronization. The third experiment characterized the frequency response by switching on/off all ACs over a range of different frequencies. The fourth experiment explored on/off switching patterns that could lead to the emergence of multiple synchronous populations of ACs. The fifth experiment explored the impact of AC heterogeneity by repeating experiments 2-4 for homogeneous and heterogeneous populations. The sixth experiment explored the impact of low/high duty cycles by repeating experiments 2-4 for ACs with high and low duty cycles. The seventh experiment explored the impact of communication delays on the degree and/or time to synchronization by repeating experiments 2-4 with deterministic and randomized delays.

Fig. 4.5 shows the active and reactive power comparison between 20 virtual ACs and the 20 experimental ACs, obtained in the first experiment. The aggregate power behavior of the virtual and experimental ACs is similar. However, as mentioned above, the effect of delayed temperature dynamics observed in the experimental testbed was significant and not adequately captured in the simulation

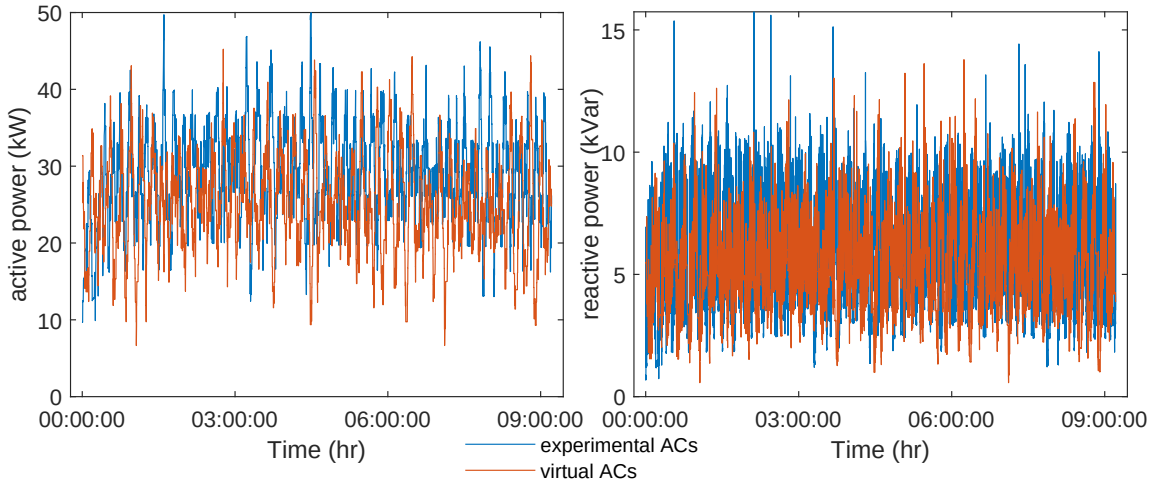


Figure 4.5: Active (left) and reactive (right) power comparison between virtual and experimental ACs.

testbed. Hence, we incorporated these dynamics into the AC simulation models. We incorporated a data-driven model leveraging the real AC data [137]; details can be found in [146].

From the other six experiments, we observed that the behavior of the experimental ACs was diverse even though the model houses were nearly identical. The heterogeneity of the experimental ACs made it nearly impossible to achieve synchronization, i.e., sustained aggregate power oscillations. Even in the second experiment, when the majority of ACs are forced into the same state by a series of switching commands, releasing the control leads to a rapid de-synchronization of the ACs. This is good news for load control; concerns about AC synchronization may be highly unlikely to materialize in practice. This effect was captured in the simulation testbed by tuning the magnitude and heterogeneity of the AC simulation model parameters. Across all experiments, we found that the simulation testbed behavior was largely in agreement with the experimental testbed behavior.

Table 4.1: Simulation Conditions

Condition	Nominal	Extreme
Reference Signal Type	PJM Reg-D	Square wave
Reference Signal Amplitude	10%	20%, 30%
Voltage Regulator Setting	Ideal	Non-Ideal
Communication Network	Perfect	Imperfect
Outdoor Temperature	$90°F$	$100°F$

4.4 Hardware-in-the-Loop Experiments

We next describe our HIL experiments, with the goals of demonstrating the capabilities and usefulness of our testbed and the ability of aggregations of ACs to provide grid services.

4.4.1 Setup

A single aggregation of both virtual and experimental ACs was coordinated to track power reference signals. To create an aggregation large enough to adequately track signals, we used 523 virtual ACs in addition to the 20 experimental ACs. The parameters of the virtual ACs were generated using GridLAB-D [118], with ±20% random variation around nominal values, and each virtual AC consumes 2.6 kW on average when on. The ACs were assigned to buses in taxonomy feeder R5-25.00-1 [144] to explore distribution network impacts.

Five conditions, shown in Table 4.1, were varied to characterize the issues that may arise when AC coordination takes place under nominal and extreme conditions. "Nominal" conditions use a historical PJM RegD signal [121] as the reference signal, signal amplitudes of 10% of the aggregation's average power consumption when not controlled (i.e., the baseline power), ideal voltage regulator operation, perfect communication network characteristics (i.e., no communication delays and no packet losses), and an outdoor temperature of $90°F$. "Extreme" conditions use a square wave reference signal (inducing a step-like response), signal amplitudes of 20% or 30%

of the baseline power, voltage regulator settings that induce over/under-voltages, imperfect communication network characteristics (i.e., delays normally distributed with mean 18 s and standard deviation 3 s, and packet loss across all input commands uniformly distributed between 5 and 10%), and an outdoor temperature of $100°F$.

It was impractical to run experiments with all possible combinations of conditions above for each controller, hence, a "Design of Experiments" approach [149] informed the experiment selection. Table 4.2 lists the 10 experimental conditions chosen and run for each controller (for a total of 30 experiments). Implementing the extreme outdoor temperature condition in the experimental testbed involved varying the heat gains in the model houses. A heat gain of 200 W approximated nominal outdoor temperature conditions, while a heat gain of 375 W was used for the extreme condition.

Table 4.2: HIL Experiment Results for the 3 controllers across the 10 experiments. Nom: nominal, Ext: extreme condition.

Case	Signal Type	Signal Amp. (%)	Voltage Reg.	Comm. Net.	Out. Temp.	NRMSE (%)			PJM Score (0-1)			Transformer overload time (s)		
						PI	Markov	PEM	PI	Markov	PEM	PI	Markov	PEM
1	Nom	Ext (20)	Nom	Nom	Nom	3.77	2.02	0.87	0.94	0.95	0.97	235	282	222
2	Ext	Ext (30)	Nom	Nom	Nom	2.76	2.26	1.03	-	-	-	316	227	257
3	Nom	Ext (30)	Nom	Ext	Nom	13.78	11.56	11.43	0.89	0.90	0.90	346	186	211
4	Nom	Nom (10)	Nom	Nom	Ext	1.67	1.39	0.57	0.94	0.95	0.96	582	502	526
5	Nom	Nom (10)	Ext	Ext	Nom	5.55	5.33	5.12	0.88	0.87	0.86	296	271	303
6	Nom	Ext (30)	Ext	Nom	Ext	4.45	2.22	0.54	0.95	0.96	0.97	979	849	1020
7	Ext	Nom (10)	Nom	Ext	Ext	2.88	2.73	3.23	-	-	-	546	696	333
8	Ext	Nom (10)	Ext	Nom	Nom	2.01	2.24	0.94	-	-	-	315	301	177
9	Nom	Ext (20)	Nom	Nom	Ext	2.82	1.67	0.55	0.95	0.96	0.97	472	504	645
10	Ext	Ext (30)	Ext	Ext	Ext	5.66	5.91	7.81	-	-	-	850	1458	784

4.4.2 Results and Discussion

Table 4.2 summarizes the experimental results across all cases. We report the Normalized Root Mean Square (NRMSE) error between the aggregate power consumption and the reference signal, normalized by the reference signal mean, for all three controllers across all ten cases. We also report the PJM performance score [131] when using the PJM signal as the reference signal, as well as the maximum consecutive overload time duration of transformers in the distribution network. We have not reported over/under-voltage violations because we have found these to be insignificant in our experiments results.

From Table 4.2 and Fig. 4.6, we see that the Markov model-based probabilistic controller and extended PEM controller have better tracking performance than the PI controller in most cases. Also, the PEM controller performs better than the Markov controller in most cases because the Markov controller uses an aggregate model to predict the power consumption of the aggregation and broadcasts probabilistic control inputs, both of which introduce some error, whereas the PEM controller uses direct-from-device turn-on/turn-off requests, which gives the aggregator a precise quantification of power flexibility, and deterministic direct-to-device inputs. The Markov controller's model-based approach provides it with an advantage over the PEM controller in imperfect communication cases 7 and 10.

Table 4.2 also shows that one of the main factors contributing to significant changes in performance was the communication network. For each controller, cases with extreme communication conditions generally had the largest tracking errors. Fig. 4.7, which shows PEM controller performance in Case 3 (imperfect communication network and a larger signal amplitude than Case 1), demonstrates the impact of communication delays and packet losses. Comparing the bottom plot of Fig. 4.6 to Fig. 4.7, we can see a clear shift of the AC power compared to the reference signal. Comparing cases 1 and 2 where all conditions except signal type and amplitude were held constant, we see that both conditions exert a more benign impact on performance compared to that of the communication condition; this also holds true for outdoor temperature when considering its effect in cases 4, 6, 8, and 9 against other

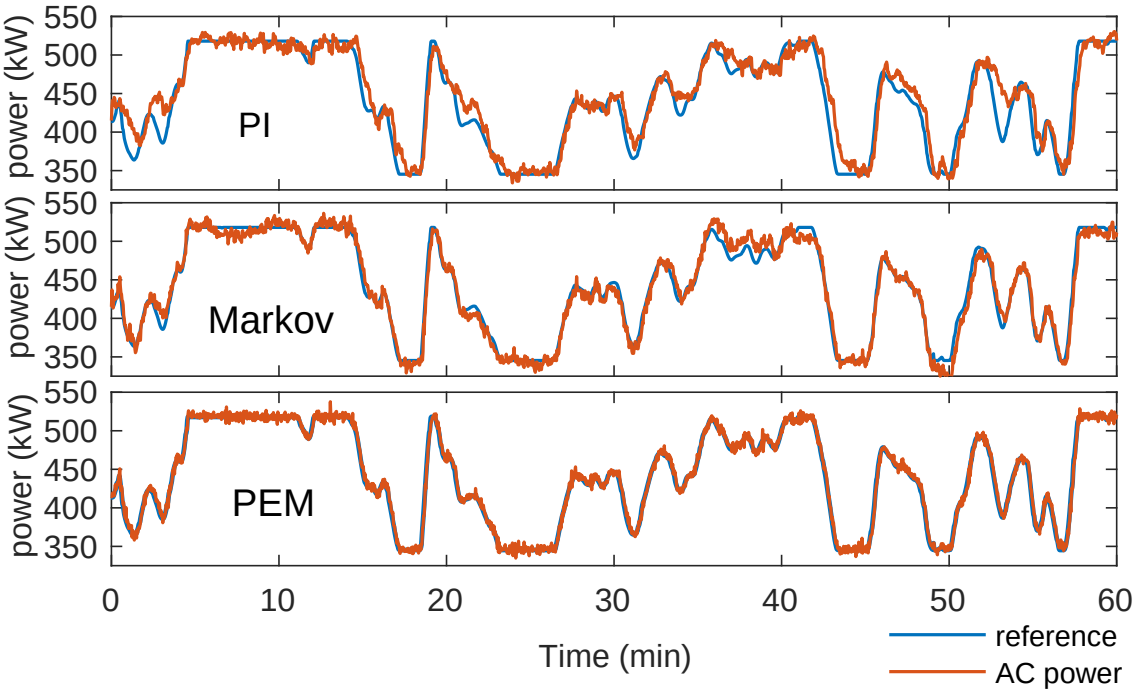


Figure 4.6: Comparison of controller tracking performance in Case 1.

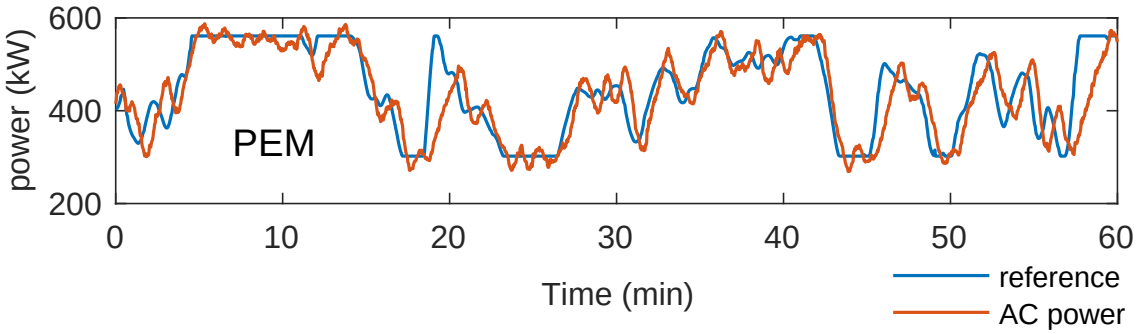


Figure 4.7: PEM controller tracking performance in Case 3 (imperfect communication network and larger amplitude reference signal than Case 1).

cases with imperfect communication. Note that in the most extreme case (Case 10), the PI controller performs the best, likely because in that case the Markov model is not very accurate and the PEM controller suffers from the imperfect communication network.

Concerning transformer overloading, we see from Table 4.2 that there is no clear trend in the maximum transformer overload time duration, leading us to conclude that all the controllers had similar effects on transformer overloading. Fig. 4.8 shows the power flowing through the transformers in Case 1 for the PEM controller. Most power flows that exceed transformer ratings (1 p.u.) are instantaneous and/or very small in magnitude. Such overloads do not usually cause actual issues to transformers, as they are not sufficient to cause overheating. In a few cases, the maximum duration is more than a few minutes; however, in these cases, the overload magnitude was very small (around 1.01 p.u.) and would be unlikely to cause any significant transformer overheating.

We next explore how the experimental ACs fared compared to the virtual ACs. Fig. 4.9 compares their lockout and on/off switching behavior in Case 1. The experimental ACs exhibit larger fluctuations due to their much smaller number, but the virtual ACs exhibit fairly similar lockout and on/off switching rates on average, indicating that the experimental ACs were not treated significantly differently than the virtual ACs. Another way to compare experimental and virtual ACs is to

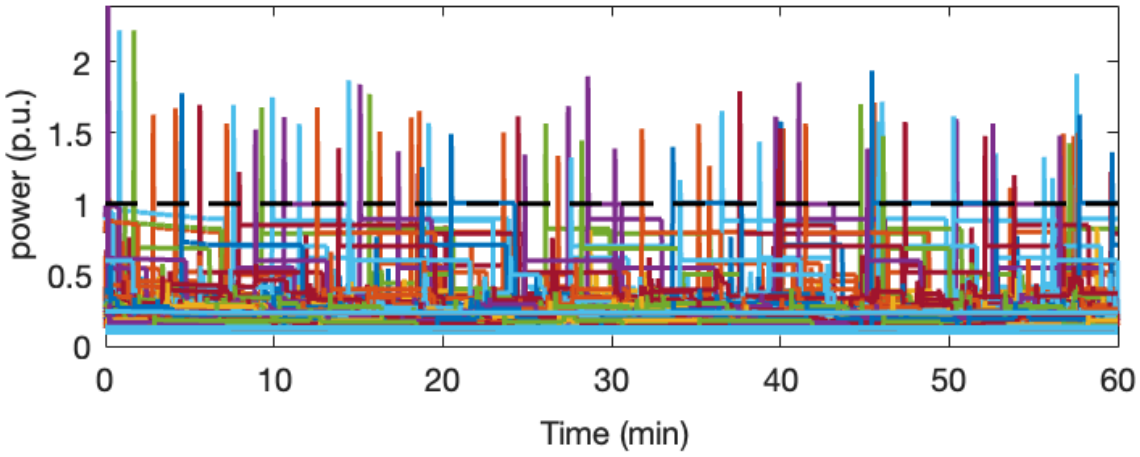


Figure 4.8: Transformer power flows in case 1 with the PEM controller. The transformer ratings were 1 p.u.

calculate the variance of the AC power from the reference signal for select groups of of 20 virtual ACs and compare the variances of those groups to the variance of the experimental ACs. We find that the experimental ACs have a variance that is generally smaller but within the range of variances of the virtual ACs, indicating that the experimental ACs are following the reference signal as well as any group of 20 virtual ACs.

Overall, from Table 4.2, we observe satisfactory tracking performances across most cases, even the extreme cases. We see the NRMSEs are less than 5% in most cases and 5-15% in a small number of cases, mainly under communication network conditions. A PJM performance score of 0.75 is required for market participation, and we see performance scores greater than 0.85 in all cases, even extreme cases, and across all controllers. Also, no oscillations, synchronization, or chaotic behavior was observed in both the extreme open-loop experiments and extreme controlled HIL experiments. These results indicate the feasibility of these load control strategies working in practical settings.

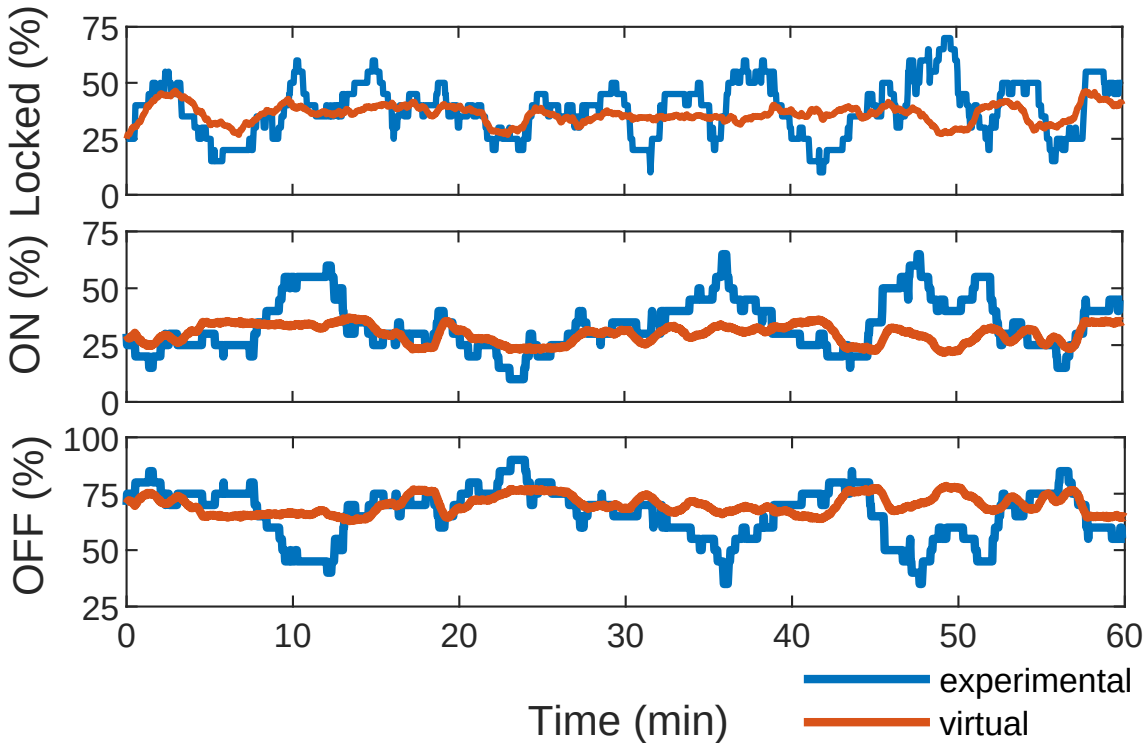


Figure 4.9: Percentage of experimental and virtual ACs that are locked (top), on (middle), and off (bottom) for the PEM controller in Case 1.

4.5 Conclusions

We have constructed a HIL experimental testbed with 20 physical model houses enabling us to test load control methods. The model houses provided data on power usage and temperature response that were used to improve high-fidelity AC simulation models. We have shown that the experimental testbed can be tuned and described how we conducted experimental testbed validation against real-world data. We performed HIL experiments to use ACs for frequency regulation leveraging three load control strategies from the literature and demonstrating that ACs are capable of providing fast timescale grid balancing services. Throughout the chapter we have highlighted the opportunities and challenges associated with testing AC load control strategies through physical experiments. In particular, experimental testing allowed us to push the system and the controller to the extreme to understand what might happen in practice under worst-case scenarios, which is usually not possible during field testing.

Future research direction could include developing delay-aware controls, because we see from our results that communication delays and packet loss have the most significant impact on performance, and the tested control strategies do not directly address this. We also acknowledge that while it was challenging to construct 20 model houses, the size remains small relative to what would be needed in practice. Hence, further testing on a larger aggregation of experimental devices might be required before such technologies could roll out at scale.

CHAPTER 5

Comparative Evaluation of Energy Modeling Methods in Residential Electrification and Decarbonization

Climate change has prioritized decarbonization initiatives for policymakers worldwide. In the United States, significant efforts are underway to decarbonize residential homes by replacing fossil fuel-dependent systems — such as space heating, water heaters, and appliances — with electric alternatives like heat pumps. Transitioning to electrified homes is crucial for achieving climate goals but requires precise evaluation methods to understand the energy, economic, and environmental impacts. This chapter broadly categorizes and assesses four energy modeling frameworks used in residential electrification studies: simple calculators, thermodynamic models, load shape modeling, and building simulation software. Focusing on model fidelity, data requirements, computational demands, and functional applications, the study highlights the strengths and limitations of each approach. Simple calculators provide quick estimates but lack precision for dynamic assessments. Thermodynamic models and building simulations offer detailed insights but demand extensive data and calibration, which could make them less suitable for long-term or large-scale studies. Load shape modeling leverages a data-driven approach that offers scalability but may lack individual household granularity and flexibility for dynamic assessments like behavioral changes in appliance usage patterns. Simplified thermodynamic models

offer a balance between data requirements and accuracy. The analysis presented here aims to guide stakeholders—-including homeowners, engineers, utility planners, researchers and policymakers—-in selecting appropriate modeling frameworks based on project complexity, objectives, and data availability. By understanding the trade-offs between accuracy and practicality, stakeholders can make informed decisions that optimize policy development and implementation strategies, facilitating the transition to a decarbonized residential sector.

5.1 Introduction

5.1.1 Aims and Background

Global sustainability and net-zero initiatives have prioritized the electrification of residential buildings as a significant step towards reducing greenhouse gas emissions [88–90]. The buildings sector contributes nearly 30% of global energy-related CO_2 emissions [88,89], with residential buildings alone accounting for approximately 17-20% [90, 150]. This substantial contribution is primarily due to thermostatically controlled loads (TCLs) – space heating, water heating, and household appliances that maintain a set temperature using thermostats – that rely on fossil fuels [135,151]. These loads are significant because they consume substantial energy and have the potential for demand flexibility, which can be harnessed for energy efficiency and grid stability. According to the IPCC's Sixth Assessment Report, rapid electrification of the residential sector is crucial to limit global warming to 1.5°C [90]. The transition to electric systems, particularly electric heat pumps for heating and cooling, offers the potential for significant reductions in carbon emissions, especially when coupled with a decarbonized electricity grid [96, 152].

With electrification initiatives on the rise, a better understanding of the decarbonization potential, environmental impacts and economic significance of the resulting electric loads is important. This requires precise evaluation models capable of assessing the impacts of electrification on energy demand, grid operations, and environmental outcomes, at both the individual home level and system at large.

The impact of home electrification can be assessed from multiple perspectives. Homeowners focus on cost, energy savings, and comfort. System operators are concerned with peak demand, demand flexibility, grid stability and infrastructure planning. Policymakers aim to meet environmental targets and promote sustainable practices. Therefore, selecting the right analytical tool depends on the specific needs and objectives of these stakeholders. Existing research predominantly emphasizes the economic and environmental outcomes of electrification, with less focus on the comparative efficacy of the different analytical methods used in these evaluations.

Various methods exist to model and analyze the effects of residential electrification, each with distinct modeling approaches, levels of complexity, data requirements, and applicability. Each method has strengths and limitations, especially in accurately estimating energy demand. From simple calculators that provide quick estimates to complex thermodynamic and building simulation models that offer detailed insights, selecting an appropriate approach is crucial for accurate assessments.

In Section 5.2, we provide a comprehensive literature review on past energy modeling comparison studies. Many prior reviews have either focused more on the results than on the methods or have been too broad in scope to consider the specifics of TCL modeling dynamics. This work aims to fill that gap by providing a systematic analysis of the relevant methods and models in the context of residential TCLs. It attempts to provide guidance for future researchers in selecting appropriate methods and models for their specific case studies and objectives in this area.

5.1.2 Motivation and Contributions

Electrification initiatives lead to an increase in flexible demand-side resources. Ref. [56] demonstrated the ancillary service potential of these loads in practice, while also refining and improving existing methods and models for understanding power and thermal dynamics of TCLs on fast timescales. Similarly, accurately quantifying the environmental and other benefits of these electric loads requires a precise understanding of their long-term (slower and longer timescales) energy dynamics. This is because environmental, economic, and other non-ancillary grid assessments (such

as future peak demand) rely on estimated energy consumption and the associated impacts of electrified loads. Hence, a closer examination of how energy dynamics are evaluated and assessed is warranted, as inaccurate model estimates could undermine decarbonization assessments and policies.

As the statistician George Box famously said, "All models are wrong, but some are useful" [153]. This work aims to highlight the usefulness and limitations of existing energy modeling approaches for residential TCLs. The high-fidelity methods and models used to qualify their ancillary service potential, as in [56], are not suitable for long-term assessments due to significant differences in timescales—seconds or minutes versus months or years. This chapter provides an overview of existing energy methods and models used in understanding and quantifying the long-term decarbonization and other non-ancillary benefits of electrification.

The scope of this work is limited to residential electric loads, particularly flexible TCLs arising from electrification, such as electric heat pumps and water heaters. We do not consider other electric loads such as electric vehicles. Residential TCLs account for over half of the residential energy consumed [135] and are the focus of this study.

This study categorizes various energy modeling methods, analyzing each by complexity, data requirements, scalability, applicability, scope, and the extent to which they balance accuracy and computational demands. We review existing literature on energy modeling methods and provide a systematic, structured comparison and application guidance across four categories, Simple Calculators, Load Shape Modeling, Thermodynamic Models and Building Simulations. We outline methods for assessing and comparing the validity, precision and accuracy of these four categories.

We describe the scenarios in which each category best applies, particularly contrasting high-fidelity models required for operational analysis with alternative methods more suited to long-term planning and hybrid frameworks that address multiple objectives.

5.1.3 Chapter Organization

The chapter is organized as follows: Section 5.2 provides a comprehensive literature review on prior studies reviewing energy modeling methods. Section 5.3 categorizes the applicable energy modeling methods into four broad categories, comparing them and describing the methodologies of each. Section 5.4 proposes a method for comparing the outlined frameworks and highlights the ensuing challenges and opportunities. Section 5.5 provides a qualitative discussion of the findings, and Section 5.6 concludes the chapter with recommendations for future research.

5.2 Literature Review

There have been numerous systematic reviews and comparative analyses of energy models in the literature, each with different scopes and focus. The energy models used for analyzing national and global energy forecasts -— along with their economic and environmental interactions -— are a separate class and diametrically different when compared to those used to study specific loads, components, or subsystems. Consequently, the reviews of these classes of models differ significantly.

5.2.1 System-Level Energy Models

Early reviews, such as those by [91] and [154], provided comprehensive overviews of energy modeling approaches and their applications to policy formation, planning, and forecasting. These studies classified energy models into several types, including sector-specific models, industry market models, integrated energy system models, and energy-economic models. They addressed the interplay between energy systems and the broader economy, considering models that employ techniques from engineering, economics, operations research, and applied mathematics. These models were instrumental in shaping energy strategies and addressing societal energy needs at national and international levels. Similarly, [94] provided a comprehensive review of system-level energy models but did not evaluate the technical features and characteristics of these models, focusing instead on comparing their outputs and final

results.

Several studies have concentrated on energy modeling within specific regions and national contexts. Ref. [155] explored how environmental costs are integrated into power system planning, focusing on methodologies used across various countries. The study evaluated analytical models and decision-making frameworks, particularly their relevance to developing countries. It highlighted challenges in incorporating environmental considerations due to constraints like limited data availability, computational demands, and geographical suitability. Similarly, [156] examined the integration of energy-economic models to address environmental challenges and policy-making in energy systems, emphasizing the need for models that capture the complexities of environmental impacts within energy planning.

Ref. [157] argued that energy policy models developed for advanced economies often fall short when applied to developing countries, failing to consider socioeconomic and infrastructural nuances. Similarly, [158] highlighted the unique challenges of modeling energy systems in developing regions, stressing the importance of incorporating local contexts into modeling efforts. Ref. [93] critically evaluated various energy system models for developing countries, focusing on their applicability to energy, environmental, and climate policy contexts. Their scope was limited to integrated models covering the energy sector and considering both supply and demand sides, excluding energy-economy interactions.

In the context of developed countries, [159] reviewed prevalent energy system models in the United Kingdom, categorizing them based on their usage and characteristics. Their review included system-level energy-economic models as well as other types of models used for energy planning and policy analysis in the United Kingdom context.

Early efforts to classify energy models were made by [92], who provided a comprehensive taxonomy of energy models based on their characteristics and applications. This work laid the foundation for understanding the diversity of energy modeling approaches and their suitability for various analytical needs. Ref. [160] reviewed advances in energy forecasting models based on engineering economics, highlighting how system-level models integrate technical and economic aspects to improve fore-

casting accuracy, particularly in industrial energy use. Ref. [161] evaluated energy efficiency policies using energy-economy models focused on the household sector, emphasizing the importance of model selection in policy evaluation. Ref. [162] investigated the gap between energy policy challenges and model capabilities, identifying limitations in existing system-level models regarding their ability to address emerging policy questions such as flexibility, sector coupling, and decentralized energy systems. Their findings underscore the need for model development to keep pace with evolving policy requirements.

While the above discussed works and others in the literature consider national and global system-level integrated energy modeling methods, there remains a gap in systematic assessments of specific energy modeling methods with a scope focused on TCLs and the US, which this work addresses.

5.2.2 Tailored Energy Models

Beyond system-level models, significant attention has been given to energy models devoted to specific loads, components, or subsystems. These reviews provide detailed insights into the methodologies, capabilities, and limitations of particular types of models or applications.

Ref. [163] conducted a detailed cross-comparison of building energy simulation tools using a reference office building as a case study. The study examined models such as EnergyPlus, Modelica, and others, focusing on their ability to accurately simulate building performance. By comparing simulation results across different tools, the authors highlighted discrepancies and areas for improvement in building simulation modeling practices. Similarly, [164] also presented a comparison of building energy simulation programs, analyzing twenty major software tools used in building energy analysis. Their work offered a comprehensive overview of the capabilities and limitations of each tool, aiding practitioners in selecting appropriate software for their needs.

Ref. [95] reviewed existing energy retrofit decision tools for homeowners, focusing on simple calculator-type models designed to assist in evaluating energy retrofit

options. The review assessed the usability, accuracy, and scope of these tools, identifying gaps and providing recommendations for enhancing their effectiveness.

Ref. [165] offered a classification and review of bottom-up energy system models, discussing their challenges and areas for improvement. Bottom-up models typically focus on detailed technological representations and are often used for modeling specific components or subsystems, such as buildings or appliances, rather than system-level energy-economic interactions.

These tailored reviews offer valuable insights into specific modeling frameworks and tools, aiding practitioners and researchers in their respective domains. However, they often do not address the modeling of TCLs in residential settings, particularly in the context of electrification and decarbonization efforts in the United States.

5.3 Energy Modeling Frameworks – Residential Electrification

The energy modeling methods used in the literature and in practice, for understanding the energy dynamics of residential electrification, can be broadly categorized into four frameworks:

1. Simple Calculators
2. Thermodynamic Models
3. Load Shape Modeling
4. Building Simulations

5.3.1 Simple Calculators

Simple calculators are essential tools in the preliminary stages of energy assessment, providing quick estimates for energy consumption and financial viability of energy-saving measures. They use basic formulas to estimate energy savings, costs and emissions associated with home electrification, utilizing default parameters like

insulation level, heating, ventilation, and air-conditioning (HVAC) efficiency, and local energy rates. These tools often employ standardized equations to provide rapid assessments based on minimal input, such as square footage, number of occupants, or average energy consumption.

Popular tools in this category include the U.S. Department of Energy's Home Energy Saver [166], Energy Star's Home Energy Yardstick [167], and various utility-provided calculators. Ref. [95] conducted a comprehensive review of existing energy retrofit decision tools for homeowners, many of which are simple calculators.

An example of a formula underpinning these tools is the basic equation for estimating annual energy consumption for heating or cooling:

$$E = \frac{HDD \times A \times U}{\eta} \times 0.024, \tag{5.1}$$

where:

- E is the annual energy consumption (kWh).
- HDD is the heating degree days ($^\circ\mathrm{C} \cdot \mathrm{day}$). Note that degree-day metrics are typically in °C (or °F), but for *temperature differences*, $1\,^\circ\mathrm{C} = 1\,\mathrm{K}$..
- A is the building envelope area (m^2).
- U is the overall heat-transfer coefficient $\big(\mathrm{W}/(\mathrm{m}^2 \cdot \mathrm{K})\big)$.
- η is the system efficiency (dimensionless).
- The factor 0.024 converts watt-days (W·day) to kilowatt-hours (kWh), since $1\,\mathrm{W} \cdot \mathrm{day} = 24\,\mathrm{Wh} = 0.024\,\mathrm{kWh}$.

The formula for energy savings that helps stakeholders understand the direct energy reduction resulting from an upgrade is:

$$\text{Energy Savings} = (E_{\text{old}} - E_{\text{new}}) \times T, \tag{5.2}$$

where:

- E_{old} and E_{new} are the *annual* energy consumptions (kWh / year) of the old and new systems, respectively.
- T is the number of years over which savings are calculated (years).

.

The payback period is another crucial financial metric calculated by:

$$\text{Payback Period} = \frac{\text{Cost of New System}}{\text{Annual Energy Savings}}, \tag{5.3}$$

where:

- Cost of New System is the upfront cost (USD).
- Annual Energy Savings is derived from the energy savings calculation adjusted for energy price or simply the yearly cost reduction in the utility bill (USD / year).

. These calculator models are useful for initial estimations where high precision is less critical. For example, tools such as Home Energy Saver require limited inputs but produce outputs that are less accurate than complex calibrated simulations [166]. While these tools lack precision, they provide essential insights at minimal computational cost. They are typically used by homeowners or real estate professionals who need preliminary estimates for energy savings or retrofit costs without detailed analysis.

Often, monthly or annual utility bills (*i.e.*, actual energy costs) are incorporated into these calculators so the results can be presented in terms of expected cost savings in dollars—making the potential benefits more tangible and relatable to the user. The U.S. Department of Energy's Energy Star savings calculator [166], for example, uses average utility rates and typical usage patterns to estimate potential savings and payback periods for various energy-efficient appliances and retrofitting projects.

While these tools do not provide highly detailed or fully accurate results, they offer a quick snapshot that helps consumers and businesses make informed preliminary decisions. For more accurate and detailed analysis, stakeholders need more

comprehensive tools that account for the complexities of energy systems and building dynamics. However, for a quick scan to determine the potential for cost savings, simple calculators serve as an invaluable resource.

In summary, simple calculators have the benefits of being user-friendly, requiring minimal technical knowledge, and providing quick estimates useful for initial assessments. However, they are limited in accuracy due to oversimplification and lack flexibility to model complex scenarios or dynamic behaviors.

5.3.2 Thermodynamic Models

Thermodynamic models simulate the physical and thermal interactions within a building to predict the performance and energy consumption of these systems under various conditions. These models consider the principles of heat transfer, fluid dynamics, and laws of thermodynamics to provide detailed analysis of energy transformations and efficiencies. Thermodynamic models can be higher fidelity, using detailed mathematical representations to estimate energy consumption based on thermal properties and they can incorporate historical (or real-time) data where available for enhanced accuracy. Space heating/cooling are often modeled using the equivalent thermal parameter (ETP) model of [45, 118–120] with indoor air temperature and inner mass temperature dynamics. The model is detailed earlier in Chapters 2.3.1 and 3.3.1. Ref. [56] employs this model to demonstrate how ensembles of TCLs can provide ancillary services like frequency regulation.

Thermodynamic models can capture detailed interactions of thermal mass and heat flow, often delivering higher accuracy than simple calculators—especially for short operational timescales. However, the practical burdens of using these models depend on how complex they are and on the parameter-estimation approach. A two-state (air temperature and on/off mode) model may be relatively simple computationally and require only moderate amounts of data, while more comprehensive multi-zone or multi-parameter models can demand extensive data (e.g., mass temperatures) and [52] shows "the parameters of such models cannot always be estimated from readily available measurements". As a result, researchers often employ sim-

plified thermodynamic representations when seeking a balance between modeling fidelity, data availability, and scalability.

For a basic single-zone home with on/off (bang-bang) control of the air conditioner or heat pump, the cooling power can only take two values: zero (off) or a fixed maximum rate Q_{rated}. When the indoor air temperature T_a exceeds the desired setpoint plus a deadband, the system switches on ($m(t) = 1$), and when it falls sufficiently below that band, the system switches off ($m(t) = 0$). The first-order differential equation then becomes:

$$C \frac{dT_a(t)}{dt} = -\frac{T_a(t) - T_o(t)}{R} \; - \; m(t)\, Q_{\text{rated}} \; + \; Q_{\text{int}}(t), \tag{5.4}$$

where:

- $T_a(t)$ is the indoor air temperature (°C),
- $T_o(t)$ is the outdoor temperature (°C),
- C is the effective thermal capacitance of the house (Wh/°C),
- R is the thermal resistance (h·°C/Wh),
- Q_{rated} is the maximum cooling power (W) when the air conditioner is running,
- $Q_{\text{int}}(t)$ is the internal heat gain from occupants and appliances (W),
- $m(t) \in \{0, 1\}$ is the binary control signal indicating off/on operation.

This approach explicitly models the discrete cycling of a typical residential air conditioner or heat pump, rather than assuming a continuously variable compressor speed. While the model is still relatively simple, it offers more realistic insight into how HVAC systems cycle in response to thermostat signals, capturing the characteristic on/off dynamics of many residential units.

This model can similarly be applied to electric water heaters, with water temperature and air temperature surrounding the tank replacing indoor and outdoor air temperatures respectively, in the thermal model above.

The model is versatile with use across different stages of project development, from preliminary screening to detailed analysis. Adaptable to various data availability scenarios, providing more accurate predictions than simple calculators but with less demand than full simulations. It gives balance between accuracy and computational efficiency. Ref. [168] uses this model in understanding what peak demand in a fully electrified New York city neighborhood looks like and the resulting impact on transformer degradation. Ref. [169] leverages this model in estimating the physical and economic impacts on distribution grids for electrification scenarios in 48 US states.

This simplified model finds use with utility planners and researchers, where a combination of macro-level assessments and detailed technical evaluations are required. They are also useful in utility planning and large-scale retrofit programs where different building types and interventions are considered.

5.3.3 Load Shape Modeling

Load shape modeling is a data-driven approach that utilizes existing load profiles of fossil-fuel-based energy consumption to estimate the corresponding electric load profiles after electrification [170]. By transforming historical fossil fuel consumption data into equivalent electrical energy demands, this method provides insights into how electrification will impact energy consumption patterns.

In this approach, the existing load shapes for heating, water heating, and other fossil-fuel-powered appliances are converted to electric load shapes by considering factors such as equipment efficiency and fuel conversion rates. The methodology involves:

1. Data Collection: Gathering historical fossil fuel consumption data (e.g., natural gas usage) with temporal resolution (hourly, daily).

2. Load Profiling: Creating load shapes that represent the temporal profile of energy consumption.

3. Conversion to Electric Loads: Applying conversion factors to transform fossil

fuel load shapes into electric load shapes, accounting for differences in equipment efficiency (e.g., using the coefficient of performance for heat pumps).

4. Adjustment Factors: Incorporating population density, behavioral changes, energy efficiency improvements, and technological advancements.

Load shape modeling is suitable for regional or utility-scale studies where aggregate impacts are of interest. However, it may lack the granularity to capture individual household variations or detailed system dynamics. This method is particularly useful for assessing the impacts of electrification on distribution grids and energy infrastructure. For example, [170] utilized load shape modeling to evaluate whether distribution grid infrastructure in Northern California could accommodate increased residential electrification and electric vehicle adoption. By modeling the projected electric loads, they identified potential infrastructure constraints and informed planning decisions. Similarly, [171] used this method in allocating electric vehicle charging demand to distribution grid circuits, considering population density.

Load shape modeling typically requires access to historical consumption data at an hourly (or finer) resolution and some form of statistical analysis to convert it into electric load shapes. While the underlying computations are relatively straightforward, assembling sufficiently granular and reliable data can be a challenge, as fuel-use records are not always publicly available or collected at a high temporal resolution. It is suitable for regional or utility-scale studies where aggregate impacts are of interest. However, it may lack the granularity to capture individual household variations or detailed system dynamics.

5.3.4 Building Simulation

Building simulation tools provide comprehensive modeling of all energy flows in a building, including thermal load calculations, system simulations, and energy consumption across different building sections and systems. These tools are designed to assist in designing, commissioning and operating high-performance buildings by simulating the impact of various technologies and architectural choices. They operate

on first-principle physics to capture building heat flows and are particularly useful in operational studies. These models require substantial data, including climatic conditions, building geometry, material properties and occupancy schedules.

Building simulation tools like BEopt [172] provide comprehensive analyses of all energy flows within a residential building, using algorithms based on the heat balance method to determine thermal loads. The energy consumption simulation is performed based on various inputs such as:

- Building geometry and orientation: Detailed 3D models of the building.
- Material properties and construction details: Thermal properties of walls, roofs, windows, etc.
- HVAC systems and controls: Specifications of heating, cooling, and ventilation systems.
- Occupant schedules and behaviors: Usage patterns affecting internal gains.
- Weather Data: Temperature, solar radiation, wind speed, and direction..

The simulations are based on solving heat balance equations for each thermal zone, accounting for conduction, convection, radiation, and mass transfer.

For example, the heat balance equation for a zone is:

$$\sum Q_{\text{cond}} + \sum Q_{\text{conv}} + \sum Q_{\text{rad}} + Q_{\text{int}} = C_{\text{zone}} \frac{dT_{\text{zone}}}{dt}, \tag{5.5}$$

where:

- Q_{cond} represents the heat flux through conduction (W).
- Q_{conv} represents the heat flux through convection (W).
- Q_{rad} represents the net radiation heat flux (W).
- Q_{int} represents internal heat gains (W) from occupants, equipment, lighting, and other sources.

- C_{zone} is the thermal capacitance (J/K or Wh/°C) of the air and other materials inside the zone.
- T_{zone} is the zone (indoor) air temperature (°C).
- t is time (s, h, or another appropriate unit).

In this equation, the left-hand side aggregates all heat gains and losses to the zone via conduction, convection, radiation, and internal sources. The right-hand side multiplies the zone's thermal capacitance by its rate of temperature change. This balance ensures that any net heat flow into the zone (positive or negative) results in a corresponding change in the zone's temperature over time.

These building simulation tools provide detailed outputs on energy consumption, system performance, and comfort metrics. They are critical for designing energy-efficient buildings and are extensively documented in the BEOpt user guide [172].

Ref. [173] uses EnergyPlus for building demand simulations across climate scenarios while assessing how climate change and heating system electrification in California may affect greenhouse gas emissions and grid capacity requirements. Similarly [174] use ResStock and EnergyPlus in examining the implications of electrifying residential space heating across Texas on the energy demand, peak loads and emissions for the ERCOT grid. Ref. [163] provided a detailed cross-comparison of building energy simulation tools highlighting their capabilities and limitations. This aids practitioners in selecting appropriate tools for their specific needs.

Building simulations offer detailed insight into whole-building energy dynamics and can simulate the impact of changes in building design, occupancy patterns, technology, and energy policies. However, they require detailed architectural, material, and behavioral data to accurately model a building [163, 164]. Calibration and validation with actual building performance data are necessary for accurate predictions. The computational resources needed can vary significantly based on model complexity and the number of simulations required. For instance, running an annual EnergyPlus simulation with hourly time steps for a single family home might take a few minutes on a standard desktop computer [175]. On the other hand, large-scale parametric studies—e.g., analyzing hundreds or thousands of different building

configurations, climates, and technology options—can easily require several hours of computing time on a high-performance workstation or cluster [176]. As a result, practitioners must balance the precision and detail of full building simulations against practical constraints in time, data availability, and computational cost.

5.4 Proposed Methodology for Comparison across Energy Modeling Frameworks

This section outlines our plan to empirically compare the four modeling frameworks described above—Simple Calculators, Thermodynamic Models, Load Shape Modeling, and Building Simulations—using real-world data from a set of residential homes provided by our collaborators Pecan Street Incorporated (PSI). Specifically, we will collect time-stamped indoor air temperature, outdoor temperature, and AC power consumption for each participating home, along with basic building characteristics such as square footage. Below, we describe (1) the data to be gathered, (2) how each modeling approach will be implemented using these data, and (3) how we will compare model outputs.

5.4.1 Data to Be Gathered from PSI Homes

PSI will equip each home with sensors and data loggers to capture:

- **Indoor temperature:** Measured at a suitable interior location (e.g., living room) or obtained from thermostat, if accessible.
- **Outdoor temperature:** From a local weather station or on-site sensor.
- **AC power consumption:** Real power draw (W or kW), time-stamped with a resolution of 5–15 minutes.
- **Home characteristics:** Square footage, approximate insulation level, and HVAC equipment rating.

Data collection will span multiple weeks or months to capture a range of weather and occupancy conditions. Each modeling framework will then use these data (along with engineering or statistical assumptions) to produce estimates of long-term energy consumption and peak demand under various scenarios.

5.4.2 Framework 1: Simple Calculators

1. Use basic building descriptors (e.g., square footage, insulation assumptions) and degree-day data from the local climate to derive monthly energy consumption. Specifically, for each month, we will apply a degree-day formula or a calculator-based tool (e.g., U.S. DOE's Home Energy Saver) to estimate the heating or cooling load, then sum across all months to generate an annual total.

2. Incorporate HVAC nameplate efficiency (or a default coefficient of performance) to convert thermal load into estimated kWh consumption.

3. Compare these monthly and annual estimates against the measured PSI data. Where we have actual monthly utility bills, we will directly compare the aggregated calculator predictions to observed consumption for the same period.

We will compare these results with the measured consumption data on a monthly or seasonal timescale, recognizing that simple calculators do not capture short-term variations or on/off cycling. This straightforward approach gives a low-effort baseline for assessing how well more advanced methods might improve upon it.

5.4.3 Framework 2: Thermodynamic Model

We next employ a simplified on/off ETP model, consistent with common methods for TCLs [31, 45, 120]. Let $T_a(k)$ be the indoor air temperature at time step k, and $m(k) \in \{0, 1\}$ the binary control signal indicating whether the air conditioner is off ($m = 0$) or on ($m = 1$). A first-order discrete-time model (adapted from [31]) is:

$$T_a(k+1) = a\,T_a(k) \;+\; b\,T_o(k) \;-\; c\,m(k) \;+\; \varepsilon(k), \tag{5.6}$$

where $T_o(k)$ is the outdoor temperature, and $\varepsilon(k)$ captures unmodeled disturbances (e.g., internal gains). The parameters a, b, c relate to the thermal resistance R, the thermal capacitance C, and the AC's rated cooling power.

Parameter Identification: To fit the model, we use historical data of $(T_a(k), T_o(k), m(k))$. Since $m(k)$ can be inferred from the measured AC power (above a certain threshold, the unit is "on"), we can apply standard linear regression or least-squares fitting techniques:

$$\begin{bmatrix} T_a(2) \\ \vdots \\ T_a(N+1) \end{bmatrix} = \begin{bmatrix} T_a(1) & T_o(1) & m(1) \\ \vdots & \vdots & \vdots \\ T_a(N) & T_o(N) & m(N) \end{bmatrix} \begin{bmatrix} a \\ b \\ -c \end{bmatrix} + \begin{bmatrix} \varepsilon(1) \\ \vdots \\ \varepsilon(N) \end{bmatrix}. \tag{5.7}$$

Literature [48,49,51] has demonstrated that such ETP models can provide reasonable accuracy with modest data requirements if the on/off operation is well-captured. We will then use the identified parameters to forecast daily or seasonal energy consumption and compare with measured AC power over the test horizon.

5.4.4 Framework 3: Load Shape Modeling

Load shape modeling uses actual consumption patterns to extrapolate how electrification will affect overall demand profiles:

1. We will aggregate the measured AC power data (at hourly or sub-hourly intervals) from PSI homes to form a representative load profile under existing conditions.

2. By adjusting for building characteristics and using known COPs or efficiency factors, we can simulate how a switch from fossil-fuel appliances to electric systems shifts the time-based load.

3. Where similar homes have been (or will be) fully electrified, we can use their real-world data to validate the model's estimated energy savings and decar-

bonization impacts. This "ground truth" comparison is especially valuable for confirming how well load-shape adjustments reflect actual system changes.

Unlike the thermodynamic or building simulation approaches, load shape modeling does not explicitly capture internal thermal dynamics. However, it can readily incorporate historical usage patterns, occupant behaviors, and local weather, making it a practical tool for forecasting aggregated or community-level impacts of electrification.

5.4.5 Framework 4: Building Simulation

Lastly, we aim to create simplified building simulation models for each home using standard software (e.g., BEopt or EnergyPlus). We will:

1. Use each home's square footage, approximate construction type, and basic occupancy assumptions.
2. Calibrate or validate the simulation by comparing short-term outputs (e.g., daily HVAC consumption) with the measured PSI data.
3. Run parametric scenarios to evaluate the long-term energy use, peak loads, or comfort metrics under different weather years or retrofit assumptions.

Building simulations can be more data-intensive than the other approaches, though we intend to keep the models relatively simple by focusing on a single-zone approximation, occupant schedules drawn from time-of-use data, and a calibrated HVAC system based on the measured runtime.

5.4.6 Comparison Metrics and Forecasting Procedure

To evaluate the performance of each approach, we will:

- Compare modeled versus measured AC consumption over the test period, examining metrics such as root mean square error (RMSE), mean absolute percentage error (MAPE), and peak demand timing.

- Extrapolate consumption forecasts to a full cooling or heating season (as applicable) to gauge how well each method predicts total kWh use and peak loads.
- Document the data requirements, run-times, and effort needed to implement each approach.

This procedure will clarify how each framework might be used in practice for planning or policy decisions: *e.g.*, whether the added complexity of a full building simulation yields significantly better accuracy compared to the simpler thermodynamic model or load shape approach.

5.4.7 Outlook for Empirical Validation

Following the data collection from PSI, we will implement each of the four methods in parallel for the same set of homes. In doing so, we will highlight practical trade-offs in complexity, data needs, and forecast accuracy. By applying the same set of performance metrics across all approaches, the results will reveal which models are most suitable for different stakeholder objectives, whether it be utility-scale planning, household-level retrofit evaluation, or longer-term policy-making. In particular, we expect:

- **Simple Calculators** to offer quick, approximate estimates, likely under-predicting or over-predicting short-term peaks.
- **Thermodynamic Models** (on/off ETP) to capture daily cycling patterns with reasonable accuracy, provided that we can reliably identify R, C and HVAC capacity from the measured data.
- **Load Shape Modeling** to deliver a direct time-based demand profile, useful for quick scenario analyses but limited in capturing occupant or weather variations beyond the historical record.
- **Building Simulations** to provide detailed insights but require additional calibration effort and greater computational overhead.

Table 5.1: Comparison of Energy Modeling Frameworks

Criteria	**Simple Calculators**	**Thermodynamic Models**	**Load Shape Modeling**	**Building Simulations**
Data Requirements	Minimal; basic building and system parameters	Moderate; thermal properties and some operational data	Moderate; historical consumption data	Extensive; detailed building geometry, materials, occupancy schedules
Skill Requirements	Low; suitable for non-experts	Moderate; requires understanding of thermodynamics	Moderate; knowledge of data analysis	High; expertise in simulation software and building science
Computing Demand	Low; quick calculations	Moderate; simulations may be computationally intensive	Moderate; depends on data size	High; simulations can be time-consuming
Scope	Preliminary energy estimates; financial viability	Detailed system dynamics and grid analysis	Assessing aggregate impacts of electrification	Comprehensive building performance evaluation
Geographical Coverage	Local	Local to regional	Regional to national	Local; to national
Time Horizons	Short-term; immediate assessments	Short to long-term	Medium to long-term planning	Short to long-term
Methodology	Basic algebraic equations	Physics-based differential equations	Data-driven statistical analysis	Physics-based simulations; numerical methods
Examples	Home Energy Saver, Energy Star Calculator, Home Energy Yardstick	ETP Models, First-Order Thermal Models	Utility Load Shapes	BEopt, ResStock

5.5 Discussion

The comparative analysis shown in Table 5.1 of the four energy modeling frameworks —- Simple Calculators, Thermodynamic Models, Load Shape Modeling, and Building Simulations -— highlights the trade-offs between complexity, data requirements, accuracy, and applicability. Each framework offers unique advantages and faces specific limitations, making them suitable for different purposes and stakeholders.

5.5.1 Comparison of Modeling Frameworks

5.5.1.1 Data and Computational Requirements

Simple Calculators require minimal data inputs and computational resources, making them accessible to non-experts for preliminary assessments. In contrast, Building Simulations demand extensive data on building geometry, materials, occupancy, and detailed weather information, along with significant computational power. Thermodynamic Models and Load Shape Modeling fall between these extremes, requiring moderate data and computational resources.

5.5.1.2 Accuracy and Detail

Building Simulations provide the highest level of detail and accuracy by modeling all energy flows within a building based on first-principle physics. However, they are highly parameterized requiring hundreds to thousands of parameters, which are often unavailable. Thermodynamic Models offer detailed insights into system behaviors, capturing the thermal dynamics of specific components like HVAC systems, using as few as one to three parameters. Load Shape Modeling provides aggregate estimates of energy consumption patterns, which are useful for regional planning but may lack granularity at the individual building level. Simple Calculators offer the least accuracy due to their reliance on simplified equations and default parameters.

5.5.1.3 Technical Expertise and Usability

The complexity of setting up and running Building Simulations requires a high level of technical expertise in building science and familiarity with simulation software. Thermodynamic Models also require a solid understanding of thermodynamics, system identification and linear system theory. Load Shape Modeling necessitates skills in data analysis and statistical methods. Simple Calculators are user-friendly and designed for stakeholders with limited technical backgrounds.

5.5.1.4 Applicability and Scope

The choice of modeling framework depends on the specific objectives and the stakeholder's needs:

- Simple Calculators are suitable for homeowners and real estate professionals seeking quick estimates of energy savings and retrofit costs.
- Thermodynamic Models are ideal for engineers and researchers focused on designing and optimizing specific systems under various conditions.
- Load Shape Modeling is beneficial for utility planners and policymakers assessing the aggregate impacts of electrification on regional or national scales.
- Building Simulations are essential for architects, engineers, and energy consultants involved in the detailed design and certification of energy-efficient buildings.

5.5.2 Implications for Stakeholders

Understanding the strengths and limitations of each modeling framework is crucial for stakeholders to select the most appropriate tool for their needs:

5.5.2.1 Homeowners and Real Estate Professionals

Simple Calculators provide accessible and quick assessments but may not capture the full potential of energy savings or account for specific building characteristics.

Encouraging the use of more detailed tools, when feasible, can lead to better-informed decisions.

5.5.2.2 Engineers and Researchers

For detailed system design and optimization, Thermodynamic Models and Building Simulations offer valuable insights. However, the selection between the two depends on the project's scope, available data, and required level of detail.

5.5.2.3 Utility Planners and Policymakers

Load Shape Modeling allows for the assessment of electrification impacts on the grid infrastructure and energy policies at a macro level. Combining this approach with more detailed models can enhance planning accuracy.

5.5.3 Balancing Accuracy and Practicality

There is a clear trade-off between model complexity and usability. While high-fidelity models like Building Simulations provide detailed results, they may not be practical for large-scale studies due to data and computational constraints. Simplified thermodynamic models offer greater scalability without sacrificing accuracy. They also find use in hybrid modeling where they are used within optimization frameworks for jointly assessing and optimizing other objectives such as user thermal comfort or grid impacts, as seen in [168, 169].

5.5.4 Trends and Challenges

As residential electrification and decarbonization efforts accelerate, the demand for accurate and scalable energy modeling tools will grow. Key challenges and trends include:

- Data Availability: Access to high-quality, granular data is essential for improving model accuracy. Initiatives to collect and share data can benefit all stakeholders.

- Model Integration: Combining different modeling frameworks can provide more comprehensive insights but requires careful handling to ensure consistency and reliability.
- User-Friendly Tools: Developing intuitive interfaces and reducing the learning curve for complex models can make advanced tools more accessible to a broader audience.

5.6 Conclusion

This chapter has provided a comprehensive review and systematic analysis of energy modeling methods applicable to residential TCLs in the United States. By categorizing and evaluating four modeling frameworks—Simple Calculators, Thermodynamic Models, Load Shape Modeling, and Building Simulations—we have highlighted their respective strengths, limitations, and suitability for different stakeholders and objectives.

Accurate energy modeling of residential TCLs is crucial for realizing the full benefits of electrification and decarbonization initiatives. By understanding the capabilities and limitations of existing modeling frameworks, stakeholders can make informed decisions that contribute to more sustainable and efficient energy systems. Continued collaboration between researchers, policymakers, industry professionals, and end-users is essential to advance modeling practices and support the transition to a low-carbon future.

CHAPTER 6

Conclusion

In this chapter, we provide a summary of insights learned from the work presented in previous chapters. We also propose future research ideas and extensions of the work done in this dissertation; they will not be addressed in this dissertation but may be explored by other researchers

6.1 Dissertation Summary and Contributions

In this work, we have demonstrated using simulations, physical experiments and real-world pilot demonstrations, the capacity for aggregations of thermostatically controlled loads (TCL) such as air-conditioners to offer ancillary services like frequency regulation. This work employed device-driven load control approaches to TCL coordination, and has shown that model-free device-driven approaches can perform at similar levels to the aggregator-driven centralized model predictive alternatives. While model-free approaches may lack the visibility and central controllability offered by aggregator-driven approaches and are often more communication-dependent, they can be easily deployed and provide similar guarantees in scenarios where the load aggregator does not have a model of its population. The operational requirements of device-driven approaches do not require models, but as illustrated in this dissertation, they can enhance operational performance. Hence, analytical models for device-driven approaches were developed in this dissertation, furthering

our understanding of dynamics associated with this approach and offering additional performance benefits to the aggregator.

While indeed this dissertation has shown TCLs offer balancing services in practice, there still exist several barriers to their large scale adoption that need to be addressed in a systemic manner as opposed to innovative workarounds. This includes more open protocol access by TCL manufacturers to power system operators as well as delay-aware load control strategies capable of handling latency affiliated with practical scenarios.

In conclusion, aggregations of distributed energy resources like TCLs are primed and tipped to play a significant role in modern power systems with increased renewable penetration.

6.2 Possible Future Research Extensions

- Delay-aware Load Control Strategies. In Chapters 3 - 4, the adverse impact of communication network delays and losses are shown to be non-trivial. It is therefore essential for load controllers to be robust to the latency and communication drop-outs.

- Larger pilot demonstrations are needed. There exist a few and limited hardware-in-the-loop testbeds for assessing and validating load control strategies, more and larger testbeds will be useful. The same is applicable to the cyber-physical test systems. There is the possibility for technical challenges not seen in the smaller testbeds (< 100 devices), as was used in this dissertation, to occur as the technology scales to larger systems. Hence, repeated testings in larger pilot programs should be undertaken before large scale adoption.

- Understanding the effects of flexible load control on load lifespan. The flexible shifting of loads involves frequent interruptions to the regular load duty-cycle as well as increased load cycling. It is necessary to understand the degradations, if any, and how severe, to the load health and life-span.

APPENDIX

Appendix A

Communication Architecture enabling Cyber-Physical Load Control for Frequency Regulation

The purpose of this appendix is to provide details on the communication architecture used in the real world demonstration of TCLs providing frequency regulation, as shown in [56]. Fig. A.1 represents the architecture of the cyber-physical system used in this demonstration.

In Fig. A.1, the master controller which executes the load control and coordination strategy being tested was implemented in MATLAB. The simulation testbed is also implemented in MATLAB. The power system operator and desired aggregate power command shown in Fig. A.1 that is used in this work is the PJM Frequency regulation signal [121], which is also implemented in MATLAB. Communication and data exchange of virtual TCL measurements such as temperature, ON/OFF mode and power between the master controller and simulation testbed is easily implemented, as both subsystems reside in MATLAB.

Fig. A.2 illustrates the communication architecture between the sub-systems in Fig. A.1. A cloud server was used for communication between the houses in the field testbed and the controllers, as seen in Fig A.1. The cloud server is capable of communicating with both the master controller in MATLAB and the physical TCLs of the field testbed. A microprocessor in the field testbed sub-system, is particu-

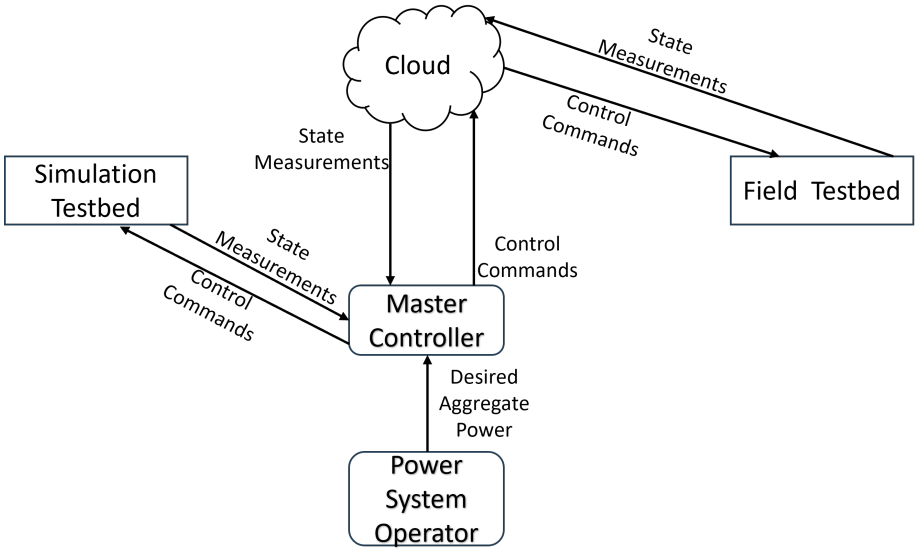


Figure A.1: Architecture of the Cyber-Physical System

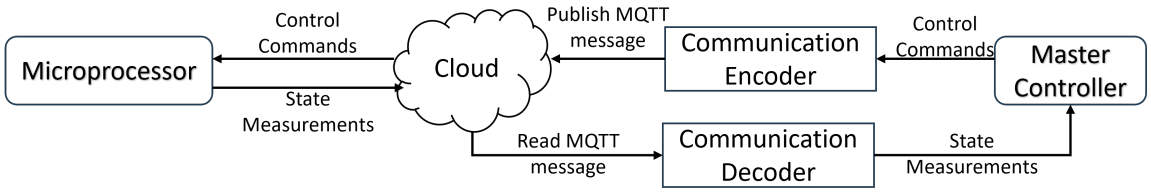


Figure A.2: Communication Architecture between Field Testbed and Controller.

larly responsible for communicating with the master controller via this cloud server. The communication and data exchange of real state measurements and control commands between the master controller and cloud server is achieved using the MQTT messaging protocol [177]. Switching control commands from the master controller are encoded as MQTT messages in MATLAB and broadcast to specific cloud server addresses (known as "topics") which the microprocessor queries often to receive instructions. Each field testbed TCL has a unique address. The microprocessors for each TCL "subscribe" to their respective topics. This is how the master controller sends varying unique control commands to each real TCL. Also, real state measurements such as temperature, ON/OFF mode, voltage and power measurements of the physical ACs are encoded by the microprocessor and published to a cloud server address for the master controller to query and collect. The master controller is able to subscribe to the "topics" of all TCLs, that is, the controller periodically queries the cloud server addresses for all TCLs for updated state measurements or can subscribe to individual "topics" for specific state measurement updates. These encoded state measurements are subsequently decoded and parsed by data filters built into the master controller to detect and flag corrupt data (infeasible or out of bound state measurements) before being passed on for computation in the master controller. The MQTT protocol was chosen for its lightweight architecture and speed, which should enable easy frequent communication between the TCL microprocessors and central controllers.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] A. V. Meier, *System Performance in "Electric Power Systems: A Conceptual Introduction"*, 2006, pp. 229–258.

[2] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2011.

[3] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl, "Demand dispatch," *IEEE Power and Energy Magazine*, vol. 8, no. 3, pp. 20–29, 2010.

[4] W. J. Hausman and J. L. Neufeld, "Time-of-day pricing in the u.s. electric power industry at the turn of the century," *The RAND Journal of Economics*, vol. 15, no. 1, pp. 116–126, 1984. [Online]. Available: http://www.jstor.org/stable/3003674

[5] M. Roozbehani, M. A. Dahleh, and S. K. Mitter, "Volatility of power grids under real-time pricing," *IEEE Transactions on Power Systems*, vol. 27, pp. 1926–1940, 2011. [Online]. Available: https://api.semanticscholar.org/CorpusID:2288193

[6] G. S. Ledva, "Learning and Control Applied to Demand Response and Electricity Distribution Networks," Ph.D. dissertation, 2019. [Online]. Available: https://deepblue.lib.umich.edu/handle/2027.42/149905

[7] L. L. Fountain and J. L. Blackburn, "Application and test of frequency relays for load shedding [includes discussion]," *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems*, vol. 73, no. 2, pp. 1660–1668, 1954.

[8] H. E. Lokay and V. Burtnyk, "Application of underfrequency relays for automatic load shedding," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-87, no. 3, pp. 776–783, 1968.

[9] B. H. Purslove Barker, "THE CENTRALIZED CONTROL OF PUBLIC LIGHTING AND OFF-PEAK LOADS BY SUPERIMPOSED RIPPLES."

[10] C. Gellings, "The concept of demand-side management for electric utilities," *Proceedings of the IEEE*, vol. 73, no. 10, pp. 1468–1470, 1985.

[11] US Department of Energy, "The American Recovery and Reinvestment Act Smart Grid Highlights: Jumpstarting a Modern Grid," no. October, p. 72, 2014.

[12] U. D. of Energy, "Benefits of demand response in electricity markets and recommendations for achieving them - a report to the united states congress pursant to section 1252 of the energy policy act of 2005," U.S. Department of Energy, Tech. Rep., 02/2006 2006.

[13] R. Hledik and K. Peters, "Real Reliability: The Value of Virtual Power VOLUME I: SUMMARY REPORT MAY 2023," vol. I, no. May, 2023.

[14] Federal Energy Regulatory Commission, "Order No. 719 — Federal Energy Regulatory Commission," https://www.ferc.gov/media/order-no-719.

[15] ——, "Order No. 745 — Federal Energy Regulatory Commission," https://www.ferc.gov/power-sales-and-markets/demand-response.

[16] J. Wellinghoff, P. D. Moeller, J. R. Norris, and C. A. Lafleur, "Order No. 755 — Federal Energy Regulatory Commission."

[17] Federal Energy Regulatory Commission, "FERC Order No. 2222: Fact Sheet," https://www.ferc.gov/media/ferc-order-no-2222-fact-sheet, 2020.

[18] ——, "FERC Order No. 2222 Explainer: Facilitating Participation in Electricity Markets by Distributed Energy Resources[i] — Federal Energy Regulatory Commission," https://www.ferc.gov/ferc-order-no-2222-explainer-facilitating-participation-electricity-markets-distributed-energy.

[19] N. Johnson, D. Nemtzow, L. P. Office, J. Paladino, and E. B. Wolfe, "Pathways to Commercial Liftoff: Virtual Power Plants," 2023. [Online]. Available: https://liftoff.energy.gov/vpp

[20] M. R. Almassalkhi and S. Kundu, "Intelligent electrification as an enabler of clean energy and decarbonization." [Online]. Available: https://doi.org/10.1007/s40518-023-00228-z

[21] U.S. Environmental Protection Agency, "Sources of greenhouse gas emissions," https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions.

[22] O. Ma, N. Alkadi, P. Cappers, P. Denholm, J. Dudley, S. Goli, M. Hummon, S. Kiliccote, J. MacDonald, N. Matson, D. Olsen, C. Rose, M. D. Sohn, M. Starke, B. Kirby, and M. O'Malley, "Demand response for ancillary services," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 1988–1995, 2013.

[23] J. S. MacDonald, P. A. Cappers, D. S. Callaway, and S. Kiliccote, "Demand response providing ancillary services a comparison of opportunities and challenges in the us wholesale markets," 2012. [Online]. Available: https://api.semanticscholar.org/CorpusID:168391450

[24] A. Keskar, D. Anderson, J. X. Johnson, I. A. Hiskens, and J. L. Mathieu, "Do commercial buildings become less efficient when they provide grid ancillary services?" *Energy Efficiency*, vol. 13, pp. 487 – 501, 2019. [Online]. Available: https://api.semanticscholar.org/CorpusID:255483997

[25] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, and D. S. Callaway, "Experimental demonstration of frequency regulation by commercial buildings—part ii: Results and performance evaluation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3224–3234, 2018.

[26] I. Beil, I. Hiskens, and S. Backhaus, "Frequency regulation from commercial building hvac demand response," *Proceedings of the IEEE*, vol. 104, no. 4, pp. 745–757, 2016.

[27] N. Motegi, M. Piette, D. Watson, S. Kiliccote, and P. Xu, "Introduction to commercial building control strategies and techniques for demand response – appendices," 01 2007.

[28] S. Koch, J. L. Mathieu, and D. S. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," 2011. [Online]. Available: https://api.semanticscholar.org/CorpusID:13225947

[29] E. Vrettos, C. Ziras, and G. Andersson, "Fast and reliable primary frequency reserves from refrigerators with decentralized stochastic control," *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 2924–2941, 2017.

[30] O. Oyefeso, G. S. Ledva, M. Almassalkhi, I. A. Hiskens, and J. L. Mathieu, "Control of aggregate air-conditioning load using packetized energy concepts," in *2022 IEEE Conference on Control Technology and Applications (CCTA)*, 2022, pp. 447–454.

[31] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

[32] E. Sortomme and M. A. El-Sharkawi, "Optimal scheduling of vehicle-to-grid energy and ancillary services," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 351–359, 2012.

[33] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2013.

[34] J. Taylor, J. Mathieu, D. Callaway, and K. Poolla, "Price and capacity competition in balancing markets with energy storage," *Energy Systems*, vol. 8, 02 2017.

[35] O. Mégel, J. L. Mathieu, and G. Andersson, "Scheduling distributed energy storage units to provide multiple services under forecast error," *International Journal of Electrical Power and Energy Systems*, vol. 72, pp. 48–57, 2015, the Special Issue for 18th Power Systems Computation Conference. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061515000939

[36] J. L. Mathieu and J. A. Taylor, "Controlling nonlinear batteries for power systems: Trading off performance and battery life," in *2016 Power Systems Computation Conference (PSCC)*, 2016, pp. 1–7.

[37] US Energy Information Administration, "U.S. Energy Consumption by Source and Sector, 2022," Tech. Rep., 2022. [Online]. Available: https://www.eia.gov/energyexplained/us-energy-facts/images/consumption-by-source-and-sector.pdf

[38] U.S. Energy Information Administration, "2020 Residential Energy Consumption Survey," https://www.eia.gov/consumption/residential/data/2020/.

[39] U.S. Environmental Protection Agency, "Electricity customers," https://www.epa.gov/energy/electricity-customers.

[40] U.S. Energy Information Administration, "Use of energy explained - energy use for transportation," https://www.eia.gov/energyexplained/use-of-energy/transportation-in-depth.php.

[41] U.S. Federal Energy Management Program, "Bidirectional charging and electric vehicles for mobile storage," https://www.energy.gov/femp/bidirectional-charging-and-electric-vehicles-mobile-storage.

[42] U.S. Energy Information Administration, "Table CE3.1 Annual household site end-use consumption in the U.S. 2015," https://www.eia.gov/consumption/residential/data/2015/cI&e/pdf/ce3.1.pdf, 2015.

[43] D. Bargiotas and J. Birdwell, "Residential air conditioner dynamic model for direct load control," *IEEE Transactions on Power Delivery*, vol. 3, no. 4, pp. 2119–2126, 1988.

[44] R. Pratt and Z. Taylor, "Development and testing of an equivalent thermal parameter model of commercial buildings from time-series end-use data," Pacific Northwest National Laboratory, 1994.

[45] R. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Princeton University, 1978.

[46] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin, "Aggregate model for heterogeneous thermostatically controlled loads with demand response," *IEEE Power and Energy Society General Meeting*, 2012.

[47] M. Gustafson, J. Baylor, and G. Epstein, "Estimating air conditioning load control effectiveness using an engineering model," *IEEE Transactions on Power Systems*, vol. 8, no. 3, pp. 972–978, 1993.

[48] S. El-Ferik, S. Hussain, and F. Al-Sunni, "Identification of physically based models of residential air-conditioners for direct load control management," in *2004 5th Asian Control Conference (IEEE Cat. No.04EX904)*, vol. 3, 2004, pp. 2079–2087 Vol.3.

[49] S. El-Ferik, S. A. Hussain, and F. M. Al-Sunni, "Identification and weather sensitivity of physically based model of residential air-conditioners for direct load control: A case study," *Energy and Buildings*, vol. 38, no. 8, pp. 997–1005, 2006. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S037877880500232X

[50] P. Bacher and H. Madsen, "Identifying suitable models for the heat dynamics of buildings," *Energy and Buildings*, vol. 43, no. 7, pp. 1511–1522, 2011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778811000491

[51] A. Pahwa and C. W. Brice, "Modeling and system identification of residential air conditioning load," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-104, no. 6, pp. 1418–1425, 1985.

[52] I. M. Granitsas, I. A. Hiskens, J. L. Mathieu, and G. S. Ledva, "Parameter identifiability and estimation of thermostatically controlled loads," in *2023 IEEE Belgrade PowerTech*, 2023, pp. 01–06.

[53] J. L. Mathieu, "Modeling, analysis, and control of demand response resources," Ph.D. dissertation, 2012, copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2023-03-04. [Online]. Available: https://proxy.lib.umich.edu/login?url=https://www.proquest.com/dissertations-theses/modeling-analysis-control-demand-response/docview/1081709057/se-2

[54] S. Koch, "Demand response methods for ancillary services and renewable energy integration in electric power systems," Ph.D. dissertation, ETH Zurich, 2013. [Online]. Available: https://doi.org/10.3929/ethz-a-009756530

[55] J. Mathieu et al., "Overcoming the technical challenges of coordinating distributed load resources at scale (final report)," OSTI.gov, Tech. Rep. 1992756, 2023.

[56] I. M. Granitsas, O. Oyefeso, S. Mock, H. Scott, I. A. Hiskens, J. L. Mathieu, and G. S. Ledva, "Controlling air conditioners for frequency regulation: A real-world example," in *Accepted to the IEEE Transactions on Smart Grid*, 2024.

[57] M. S. Nazir, S. C. Ross, J. L. Mathieu, and I. A. Hiskens, "Performance limits of thermostatically controlled loads under probabilistic switching," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 8873–8880, 2017, 20th IFAC World Congress. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896317321316

[58] M. S. Nazir and I. A. Hiskens, "Noise and Parameter Heterogeneity in Aggregate Models of Thermostatically Controlled Loads," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 8888–8894, jul 2017.

[59] L. C. Totu, R. Wisniewski, and J. Leth, "Demand response of a tcl population using switching-rate actuation," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 5, pp. 1537–1551, 2017.

[60] S. J. Crocker and J. L. Mathieu, "Adaptive state estimation and control of thermostatic loads for real-time energy balancing," in *2016 American Control Conference (ACC)*, 2016, pp. 3557–3563.

[61] M. Almassalkhi, J. Frolik, and P. Hines, "Packetized energy management: Asynchronous and anonymous coordination of thermostatically controlled loads," *American Control Conference*, 2017.

[62] M. Almassalkhi, L. Duffaut-Espinosa, P. Hines, J. Frolik, S. Paudyal, and M. Amini, "Asynchronous coordination of distributed energy resources with packetized energy management," in *Energy Markets and Responsive Grids*, S. Meyn, T. Samad, I. Hiskens and J. Stoustrup, Ed. Springer, 2018.

[63] L. A. Espinosa, M. Almassalkhi, P. Hines, S. Heydari, and J. Frolik, "Towards a macromodel for Packetized Energy Management of resistive water heaters," *2017 51st Annual Conference on Information Sciences and Systems, CISS 2017*, 2017.

[64] 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, 2009.

[65] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, "Modeling and control of thermostatically controlled loads," 01 2011.

[66] S. Bashash and H. K. Fathy, "Modeling and control of aggregate air conditioning loads for robust renewable power management," *IEEE Trans Control Systems Technology*, vol. 21, no. 4, pp. 1318–1327, 2013.

[67] C. Perfumo, E. Kofman, J. H. Braslavsky, and J. K. Ward, "Load management: Model-based control of aggregate power for populations of thermostatically controlled loads," *Energy Conversion and Management*, vol. 55, pp. 36–48, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0196890411002998

[68] S. Kundu and I. A. Hiskens, "Nonlinear dynamics of hysteresis-based load controls," *IFAC Proceedings Volumes*, vol. 47, no. 3,

pp. 5419–5425, 2014, 19th IFAC World Congress. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667016424572

[69] A. Lesage-Landry and J. A. Taylor, "Setpoint tracking with partially observed loads," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5615–5627, 2018.

[70] A. Lesage-Landry and D. S. Callaway, "Dynamic and distributed online convex optimization for demand response of commercial buildings," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 632–637, 2020.

[71] A. Lesage-Landry, I. Shames, and J. A. Taylor, "Predictive online convex optimization," *Automatica*, vol. 113, p. 108771, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S000510981930634X

[72] A. Lesage-Landry, H. Wang, I. Shames, P. Mancarella, and J. A. Taylor, "Online convex optimization of multi-energy building-to-grid ancillary services," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2416–2431, 2020.

[73] A. Lesage-Landry, J. A. Taylor, and D. S. Callaway, "Online convex optimization with binary constraints," *IEEE Transactions on Automatic Control*, vol. 66, no. 12, pp. 6164–6170, 2021.

[74] A. Lesage-Landry and J. A. Taylor, "Online convex optimization for demand response."

[75] A. Lesage-Landry and D. S. Callaway, "Batch reinforcement learning for network-safe demand response in unknown electric grids," *Electric Power Systems Research*, vol. 212, p. 108375, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378779622005363

[76] A. Lesage-Landry and J. A. Taylor, "Learning to Shift Thermostatically Controlled Loads," 2017. [Online]. Available: http://hdl.handle.net/10125/41522

[77] V. Mai, T. Zhang, and A. Lesage-Landry, "Multi-agent reinforcement learning for renewable integration in the electric power grid."

[78] A. D. Domínguez-García and C. N. Hadjicostis, "Distributed algorithms for control of demand response and distributed energy resources," in *2011 50th*

IEEE Conference on Decision and Control and European Control Conference, 2011, pp. 27–32.

[79] A. D. Dominguez-Garcia, S. T. Cady, and C. N. Hadjicostis, "Decentralized optimal dispatch of distributed energy resources," *IEEE Conference on Decision and Control*, pp. 3688–3693, 2012. [Online]. Available: https://experts.illinois.edu/en/publications/decentralized-optimal-dispatch-of-distributed-energy-resources

[80] X. Zhou, E. Dall'Anese, L. Chen, and A. Simonetto, "An incentive-based online optimization framework for distribution grids," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 2019–2031, 2018.

[81] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, 2014.

[82] I. Krstic M., Kanellakopoulos and P. Kokotovic, 1995.

[83] M. Krstić, "Performance improvement and limitations in extremum seeking control," *Systems I& Control Letters*, vol. 39, no. 5, pp. 313–326, 2000. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167691199001115

[84] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109899001831

[85] D. Dochain, M. Perrier, and M. Guay, "Extremum seeking control and its application to process and reaction systems: A survey," *Mathematics and Computers in Simulation*, vol. 82, no. 3, pp. 369–380, 2011, 6th Vienna International Conference on Mathematical Modelling. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378475410003290

[86] M. Almassalkhi, J. Frolik, and P. Hines, "How to prevent blackouts by packetizing the power grid - ieee spectrum," 1 2022. [Online]. Available: https://spectrum.ieee.org/packetized-power-grid

[87] A. Khurram, M. Amini, L. A. D. Espinosa, P. D. H. Hines, and M. Almassalkhi, "Real-time Grid and DER Co-simulation Platform for Validating Large-scale DER Control Schemes," pp. 1–8, 2021. [Online]. Available: http://arxiv.org/abs/2106.01993

[88] IEA, "Global energy review: Co2 emissions in 2021," International Energy Agency, Tech. Rep., 2022. [Online]. Available: https://www.iea.org/reports/global-energy-review-co2-emissions-in-2021-2

[89] ——, "World energy outlook 2022," International Energy Agency, Tech. Rep., 2022. [Online]. Available: https://www.iea.org/reports/world-energy-outlook-2022

[90] IPCC, "Climate change 2022: Mitigation of climate change," Intergovernmental Panel on Climate Change, Tech. Rep., 2022. [Online]. Available: https://www.ipcc.ch/report/ar6/wg3/

[91] K. Hoffman and D. O. Wood, "Energy system modelling and forecasting," *Annual Review of Energy*, vol. 1, pp. 423–453, 1976.

[92] N. Van Beeck, "Classification of energy models," Tilburg University and Eindhoven University of Technology, Tech. Rep., 1999.

[93] S. C. Bhattacharyya and G. R. Timilsina, "A review of energy system models," *International Journal of Energy Sector Management*, vol. 4, no. 4, pp. 494–518, 2010.

[94] M. G. Prina, B. Nastasi, D. Groppi, S. Misconel, D. A. Garcia, and W. Sparber, "Comparison methods of energy system frameworks, models and scenario results," *Renewable and Sustainable Energy Reviews*, vol. 167, p. 112719, 2022.

[95] M. Seddiki, A. Bennadji, R. Laing, D. Gray, and J. M. Alabid, "Review of existing energy retrofit decision tools for homeowners," *Energies*, vol. 12, no. 5, p. 948, 2019.

[96] S. Ebrahimi, M. Mac Kinnon, and J. Brouwer, "California end-use electrification impacts on carbon neutrality and clean air," *Applied Energy*, vol. 213, pp. 435–449, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626191830059X

[97] O. E. Oyefeso, D. A. Geller, I. M. Granitsas, D. S. Callaway, and J. L. Mathieu, "A hardware-in-the-loop experimental testbed using air conditioners for grid balancing," in *in review*, 2024.

[98] G. S. Ledva, I. M. Granitsas, O. Oyefeso, S. A. Nugroho, P. Phanivong, I. A. Hiskens, and J. L. Mathieu, "Co-simulation testbed and controllers

for compressor-based loads providing grid balancing services," 2023, GitHub repository. [Online]. Available: https://github.com/Giannisgran/Co-Simulation-testbed-and-controllers-for-compressor-based-loads-providing-grid-balancing-services

[99] J. A. Taylor, S. V. Dhople, and D. S. Callaway, "Power systems without fuel," *Renewable and Sustainable Energy Reviews*, vol. 57, pp. 1322–1336, 2016.

[100] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox, "Homeostatic utility control," *IEEE Trans Power Apparatus and Systems*, vol. PAS-99, no. 3, pp. 1151–1163, May 1980.

[101] Y. Ding, Y. Song, H. Hui, and C. Shao, *Integration of Air Conditioning and Heating into Modern Power Systems*. Springer, 2019.

[102] A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya, "Aggregate flexibility of a collection of loads," *IEEE Conference on Decision and Control*, pp. 5600–5607, 2013.

[103] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Trans Power Systems*, vol. 30, no. 1, pp. 189–198, 2015.

[104] L. W. Davis and P. J. Gertler, "Contribution of air conditioning adoption to future energy use under global warming," *Proceedings of the National Academy of Sciences*, vol. 112, no. 19, pp. 5962–5967, 2015.

[105] M. Izquierdo, A. Moreno-Rodríguez, A. González-Gil, and N. García-Hernando, "Air conditioning in the region of Madrid, Spain: An approach to electricity consumption, economics and CO2 emissions," *Energy*, vol. 36, no. 3, pp. 1630–1639, 2011.

[106] C. Perfumo, J. H. Braslavsky, and J. K. Ward, "Model-based estimation of energy savings in load control events for thermostatically controlled loads," *IEEE Trans Smart Grid*, vol. 5, no. 3, pp. 1410–1420, 2014.

[107] S. Grammatico, B. Gentile, F. Parise, and J. Lygeros, "A mean field control approach for demand side management of large populations of thermostatically controlled loads," in *2015 European Control Conference (ECC)*, 2015, pp. 3548–3553.

[108] D. Wu, J. Lian, Y. Sun, T. Yang, and J. Hansen, "Hierarchical control framework for integrated coordination between distributed energy resources and demand response," *Electric Power Systems Research*, vol. 150, pp. 45–54, sep 2017.

[109] B. Chen, J. Francis, M. Pritoni, S. Kar, and M. Bergés, "COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility," *BuildSys - ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 31–40, nov 2020. [Online]. Available: https://doi.org/10.1145/3408308.3427980

[110] M. Franceschelli, A. Pilloni, and A. Gasparri, "Multi-Agent Coordination of Thermostatically Controlled Loads by Smart Power Sockets for Electric Demand Side Management," *IEEE Trans Control Systems Technology*, vol. 29, no. 2, pp. 731–743, 2021.

[111] K. Kircher, A. Aderibole, L. K. Norford, and S. Leeb, "Distributed Peak Shaving for Small Aggregations of Cyclic Loads," *IEEE Trans Power Delivery*, vol. 8977, no. c, pp. 1–1, 2022.

[112] L. A. Espinosa, M. Almassalkhi, P. Hines, and J. Frolik, "Aggregate modeling and coordination of diverse energy resources under packetized energy management," *IEEE Conference on Decision and Control*, pp. 1394–1400, 2018.

[113] ——, "System properties of packetized energy management for aggregated diverse resources," *Power Systems Computation Conference*, 2018.

[114] L. A. D. Espinosa and M. Almassalkhi, "A Packetized Energy Management Macromodel With Quality of Service Guarantees for Demand-Side Resources," *IEEE Trans Power Systems*, vol. 35, no. 5, pp. 3660–3670, 2020.

[115] L. A. Espinosa, A. Khurram, and M. Almassalkhi, "Reference-Tracking Control Policies for Packetized Coordination of Heterogeneous DER Populations," *IEEE Trans Control Systems Technology*, 2020.

[116] H. Mavalizadeh, L. A. Duffaut Espinosa, and M. R. Almassalkhi, "Decentralized frequency control using packet-based energy coordination," *IEEE SmartGridComm*, 2020.

[117] C. Y. Chang, W. Zhang, J. Lian, and K. Kalsi, "Modeling and control of aggregated air conditioning loads under realistic conditions," *Innovative Smart Grid Technologies Conference*, 2013.

[118] "Tech:Residential – GridLAB-D Wiki," http://gridlab-d.shoutwiki.com/wiki/Tech:Residential.

[119] K. Subbarao, "Thermal parameters for single and multizone buildings and their determination from performance data," Solar Energy Research Institute, 1981.

[120] N. Wilson, B. Wagner, and W. Colborne, "Equivalent thermal parameters for an occupied gas-heated house," *ASHRAE Transactions*, vol. 91, p. 2, 1985.

[121] PJM Interconnection, "PJM - Ancillary Services," https://pjm.com/markets-and-operations/ancillary-services.

[122] CPowerEnergyManagement, "How the PJM interconnection frequency regulation program works," https://cpowerenergymanagement.com/wp-content/uploads/2019/03/PJM_FR_Snapshot_3-1-2019.pdf, p. 3, 2019.

[123] A. Khurram, L. Duffaut Espinosa, and M. Almassalkhi, "A Methodology for Quantifying Flexibility in a fleet of Diverse DERs," in *IEEE PES PowerTech*, 2021.

[124] M. Banaei, F. D'Ettorre, R. Ebrahimy, M. R. Almassalkhi, and H. Madsen, "Procuring flexibility in power systems with incentive-based grid access requests," *International Journal of Electrical Power & Energy Systems*, vol. 156, p. 109745, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0142061523008025

[125] R. T. Ruminsky, "Compressor minimum off-time system," Dec. 5 1978, uS Patent 4,128,854.

[126] B. Zhang and J. Baillieul, "A packetized direct load control mechanism for demand side management," *IEEE Conf on Decision & Control*, 2012.

[127] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[128] S. H. Tindemans, V. Trovato, and G. Strbac, "Decentralized control of thermostatic loads for flexible demand response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, 2015.

[129] Y. Wan, C. Long, R. Deng, G. Wen, X. Yu, and T. Huang, "Distributed event-based control for thermostatically controlled loads under hybrid cyber attacks," *IEEE Transactions on Cybernetics*, vol. 51, no. 11, pp. 5314–5327, 2021.

[130] S. C. Ross and J. L. Mathieu, "Strategies for Network-Safe Load Control with a Third-Party Aggregator and a Distribution Operator," *IEEE Trans Power Systems*, vol. 8950, no. c, pp. 1–11, 2021.

[131] I. Dropkin, "PJM performance score," 2022. [Online]. Available: https://www.pjm.com/-/media/committees-groups/task-forces/rmdstf/2022/20220426/20220426-item-04-performance-score-education.ashx

[132] M. Amini, A. Khurram, A. Klem, M. Almassalkhi, and P. D. Hines, "A model-predictive control method for coordinating virtual power plants and packetized resources, with hardware-in-the-loop validation," in *IEEE Power & Energy Society General Meeting*, 2019, pp. 1–5.

[133] S. Brahma, A. Khurram, H. Ossareh, and M. Almassalkhi, "Optimal Frequency Regulation Using Packetized Energy Management," *IEEE Transactions on Smart Grid*, vol. 14, no. 1, pp. 341–353, 2023.

[134] M. Kenney, J. Wahlgren, K. Duloglo, T. Mateo, D. Drozdowicz, and S. Bailey, "Final 2021 integrated energy policy report, volume i: Building decarbonization," California Energy Commission, Tech. Rep. CEC-100-2021-001-V1, Feb. 2022.

[135] U.S. Energy Information Administration. (2020) Residential energy consumption survey (RECS). [Online]. Available: https://www.eia.gov/consumption/residential/data/2020/

[136] D. 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, pp. 1389–1400, 2009.

[137] Pecan Street, Inc. (2020) Dataport. [Online]. Available: http://www.pecanstreet.org/dataport

[138] V. Lakshmanan, M. Marinelli, J. Hu, and H. W. Bindner, "Provision of secondary frequency control via demand response activation on thermostatically controlled loads: Solutions and experiences from Denmark," *Applied Energy*, vol. 173, pp. 470–480, 2016.

[139] I. M. Saleh, A. Postnikov, C. Arsene, A. C. Zolotas, C. Bingham, R. Bickerton, and S. Pearson, "Impact of demand side response on a commercial retail refrigeration system," *Energies*, vol. 11, no. 2, 2018.

[140] E. Vrettos, E. Kara, J. MacDonald, G. Andersson, and D. Callaway, "Experimental demonstration of frequency regulation by commercial buildings – Part I: modeling and hierarchical control design," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3213–3223, 2018.

[141] ——, "Experimental demonstration of frequency regulation by commercial buildings – Part II: results and performance evaluation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3224–3234, 2018.

[142] Lawrence Berkeley National Laboratory. (2023) FLEXLAB. [Online]. Available: https://flexlab.lbl.gov/

[143] DOE-IEDO, "Variable-speed, low-cost motor for residential HVAC systems," Mar 2024. [Online]. Available: https://https://www.energy.gov/eere/iedo/variable-speed-low-cost-motor-residential-hvac-systems

[144] K. P. Schneider, Y. Chen, D. P. Chassin, R. G. Pratt, D. W. Engel, and S. E. Thompson, "Modern grid initiative distribution taxonomy final report," OSTI.gov, Tech. Rep. 1040684, 2008.

[145] F. Golnaraghi and B. C. Kuo, *Automatic Control Systems*, 10th ed. New York: McGraw-Hill Education, 2017.

[146] I. Granitsas, I. Hiskens, and J. Mathieu, "Incorporating delayed temperature dynamics into residential air conditioner models for grid frequency regulation," in *IEEE Conf Control Tech & App*, 2022.

[147] D. Geller et al., "Evidence of completion of milestone 3: Experimental testbed," OSTI.gov, Tech. Rep. 1804326, 2021.

[148] ——, "Evidence of completion of milestone 4: Simulation testbed validated with experimental data," OSTI.gov, Tech. Rep. 1836964, 2021.

[149] J. C. Spall, "Factorial design for efficient experimentation," *IEEE Control Systems Magazine*, vol. 30, no. 5, pp. 38–53, 2010.

[150] UNEP, "2022 global status report for buildings and construction," United Nations Environment Programme, Tech. Rep., 2022. [Online]. Available: https://globalabc.org/resources/publications/2022-global-status-report-buildings-and-construction

[151] EIA, "Annual energy outlook 2023," U.S. Energy Information Administration, Tech. Rep., 2023. [Online]. Available: https://www.eia.gov/outlooks/aeo/

[152] NREL, "Electrification futures study," National Renewable Energy Laboratory, Tech. Rep., 2022. [Online]. Available: https://www.nrel.gov/analysis/electrification-futures.html

[153] G. E. P. Box, "Science and statistics," *Journal of the American Statistical Association*, vol. 71, no. 356, pp. 791–799, 1976.

[154] F. Wirl and E. Szirucsek, "Energy modelling – a survey of related topics," *OPEC Review*, vol. 14, no. 4, pp. 361–378, Autumn 1990.

[155] A. Markandya, "Environmental costs and power systems planning," *Utilities Policy*, vol. 1, no. 1, pp. 13–27, 1990.

[156] T. Nakata, "Energy-economic models and the environment," *Progress in Energy and Combustion Science*, vol. 30, no. 4, pp. 417–475, 2004.

[157] R. Pandey, "Energy policy modelling: agenda for developing countries," *Energy Policy*, vol. 30, no. 2, pp. 97–106, 2002.

[158] F. Urban, R. M. J. Benders, and H. C. Moll, "Modelling energy systems for developing countries," *Energy Policy*, vol. 35, no. 6, pp. 3473–3482, 2007.

[159] L. M. H. Hall and A. R. Buckley, "A review of energy systems models in the uk: Prevalent usage and categorisation," *Applied Energy*, vol. 169, pp. 607–628, 2016.

[160] E. Worrell, S. Ramesohl, and G. Boyd, "Advances in energy forecasting models based on engineering economics," *Annual Review of Environment and Resources*, vol. 29, pp. 345–381, 2004.

[161] L. Mundaca, L. Neij, E. Worrell, and M. McNeil, "Evaluating energy efficiency policies with energy-economy models," *Annual Review of Environment and Resources*, vol. 35, pp. 305–344, 2010.

[162] G. Savvidis, K. Siala, C. Weissbart, L. Schmidt, F. Borggrefe, S. Kumar *et al.*, "The gap between energy policy challenges and model capabilities," *Energy Policy*, vol. 125, pp. 503–520, 2019.

[163] M. Magni, F. Ochs, S. de Vries, A. Maccarini, and F. Sigg, "Detailed cross comparison of building energy simulation tools results using a reference office building as a case study," *Energy and Buildings*, vol. 145, pp. 87–99, 2017.

[164] D. B. Crawley, J. W. Hand, M. Kummert, and B. T. Griffith, "Contrasting the capabilities of building energy performance simulation programs," *Building and Environment*, vol. 43, no. 4, pp. 661–673, 2008.

[165] M. G. Prina, G. Manzolini, D. Moser, B. Nastasi, and W. Sparber, "Classification and challenges of bottom-up energy system models—a review," *Renewable and Sustainable Energy Reviews*, vol. 129, p. 109917, 2020.

[166] U.S. Department of Energy Home Energy Saver. Estimating appliance and home electronic energy use. [Online]. Available: https://www.energy.gov/energysaver/estimating-appliance-and-home-electronic-energy-use

[167] "Energy star home energy yardstick," https://www.energystar.gov/campaign/home-energy-yardstick, accessed:2024-11-19.

[168] P. Priyadarshan, E. Pergantis, C. Crozier, K. Baker, and K. Kircher, "Edgie: A simulation test-bed for investigating the impacts of building and vehicle electrification on distribution grids," *Proceedings of the 57th Hawaii International Conference on System Sciences*, pp. 4838–4847, 1 2024. [Online]. Available: https://hdl.handle.net/10125/106760

[169] Priyadarshan, C. Crozier, K. Baker, and K. Kircher, "Distribution grids may be a barrier to residential electrification," 2024. [Online]. Available: https://arxiv.org/abs/2410.04540

[170] S. Elmallah, A. M. Brockway, and D. Callaway, "Can distribution grid infrastructure accommodate residential electrification and electric vehicle adoption in northern california?" *Electric Power Systems Research*, vol. 189, p. 106694, 2020.

[171] C. Hecht, S. Das, C. Bussar, and D. U. Sauer, "Representative, empirical, real-world charging station usage characteristics and data in germany," *eTransportation*, vol. 6, p. 100079, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2590116820300064

[172] NREL, "Beopt™ software for building energy optimization: Features and capabilities," National Renewable Energy Laboratory, Tech. Rep., 2006.

[173] B. Tarroja, F. Chiang, A. AghaKouchak, S. Samuelsen, S. V. Raghavan, M. Wei, K. Sun, and T. Hong, "Translating climate change and heating system electrification impacts on building energy use to future greenhouse gas emissions and electric grid capacity requirements in california," *Applied Energy*, vol. 225, pp. 522–534, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261918307619

[174] P. R. White, J. D. Rhodes, E. J. Wilson, and M. E. Webber, "Quantifying the impact of residential space heating electrification on the texas electric grid," *Applied Energy*, vol. 298, p. 117113, 2021. [Online]. Available: https://doi.org/10.1016/j.apenergy.2021.117113

[175] T. Hong, F. Buhl, P. Haves *et al.*, "Comparing computer run time of building simulation programs," *Building Simulation*, vol. 1, pp. 210–213, 2008.

[176] Y. Zhang and I. Korolija, "Performing complex parametric simulations with jeplus." Shanghai Jiao Tong University, 2010.

[177] S. Quincozes, T. Emilio, and J. Kazienko, "MQTT Protocol: Fundamentals, tools and future directions," *IEEE Latin America Transactions*, vol. 17, no. 09, pp. 1439–1448, 2019.