Spring 1-1-2014

Building-To-Grid Integration Through Commercial Building Portfolios Participating in Energy and Frequency Regulation Markets

Gregory S. Pavlak
University of Colorado at Boulder, greg.pavlak@gmail.com

Building-to-Grid Integration through Commercial Building Portfolios Participating in Energy and Frequency Regulation Markets

by

Gregory S. Pavlak

B.S., Hope College, 2009

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Civil, Environmental, and Architectural Engineering

2014

This thesis entitled:
Building-to-Grid Integration through Commercial Building Portfolios Participating in Energy and Frequency Regulation Markets
written by Gregory S. Pavlak
has been approved for the Department of Civil, Environmental, and Architectural Engineering

Gregor P. Henze, Ph.D., P.E.

Michael J. Brandemuehl, Ph.D., P.E.

Moncef Krarti, Ph.D., P.E.

Rajagopalan Balaji, Ph.D.

Clemens Felsmann, Dr.-Ing.

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Pavlak, Gregory S. (Ph.D.)

Building-to-Grid Integration through Commercial Building Portfolios Participating in Energy and Frequency Regulation Markets

Thesis directed by Prof. Gregor P. Henze, Ph.D., P.E.

Building energy use is a significant contributing factor to growing worldwide energy demands. In pursuit of a sustainable energy future, commercial building operations must be intelligently integrated with the electric system to increase efficiency and enable renewable generation. Toward this end, a model-based methodology was developed to estimate the capability of commercial buildings to participate in frequency regulation ancillary service markets. This methodology was integrated into a supervisory model predictive controller to optimize building operation in consideration of energy prices, demand charges, and ancillary service revenue. The supervisory control problem was extended to building portfolios to evaluate opportunities for synergistic effect among multiple, centrally-optimized buildings. Simulation studies performed showed that the multi-market optimization was able to determine appropriate opportunities for buildings to provide frequency regulation. Total savings were increased by up to thirteen percentage points, depending on the simulation case. Furthermore, optimizing buildings as a portfolio achieved up to seven additional percentage points of savings, depending on the case. Enhanced energy and cost savings opportunities were observed by taking the novel perspective of optimizing building portfolios in multiple grid markets, motivating future pursuits of advanced control paradigms that enable a more intelligent electric grid.

Dedication

To Jami, Mom, Dad, Jeff, Janice, and Nick.

If you can dream—and not make dreams your master;

If you can think—and not make thoughts your aim;

"If" $\sim$ Rudyard Kipling

Acknowledgements

The journey toward this dissertation was neither smooth, nor straight. Gregor Henze's enthusiasm and dedication were among the few constants along the way that kept me inspired and in pursuit. It has been a privilege to work under his advise, and I am forever grateful. I would also like to thank Vincent J. Cushing, Jr. for his invaluable contributions to this research. His passion for this topic is genuine, and much inspiration and direction were gathered through discussions with him. I am also sincerely grateful for the sponsorship from Clean Urban Energy, Inc. and $\dot{Q}$Coefficient, Inc. over the years that has enabled this research. Sandro Plamp and my fellow grad students have also been indispensable resources, to whom I owe many thanks as well. Finally, the steadfast love and support from my family has been paramount in this endeavor. I simply cannot thank you all enough.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction and Motivation

1.1 Global Energy and the Case for Efficiency

It is incomprehensible to imagine a world without energy. Current figures estimate that the world consumes roughly 500 quadrillion Btus of primary energy annually to support human existence and development [23]. At the most basic level, energy enables survival by aiding in the provision of life necessities such as food, water, and habitable environments. On another level, energy is a vital component to human progress and empowers a creative, prosperous existence. The undeniable influence of energy in our lives has sparked much inquiry into the the interconnection of energy and humanity. Specifically, researchers have sought to understand whether energy access and consumption are reactants driving human development, products of societal advancement, or components of a more complex symbiosis.

To this end it is necessary to define some measure of human development. Gross domestic product (GDP) has often been used as a proxy for growth and has been found to be correlated with energy consumption [96, 103]. Correlations between GDP and energy use per capita tend to exhibit a positive relationship with GDP increasing as energy consumption increases. Alternatively, the United Nations Development Programme (UNDP) has developed the Human Development Index (HDI) in effort of providing an aggregate descriptor for a countries' state of development. The HDI quantifies development by combining information along three primary dimensions: a long and healthy life, knowledge, and a decent standard of living [69].

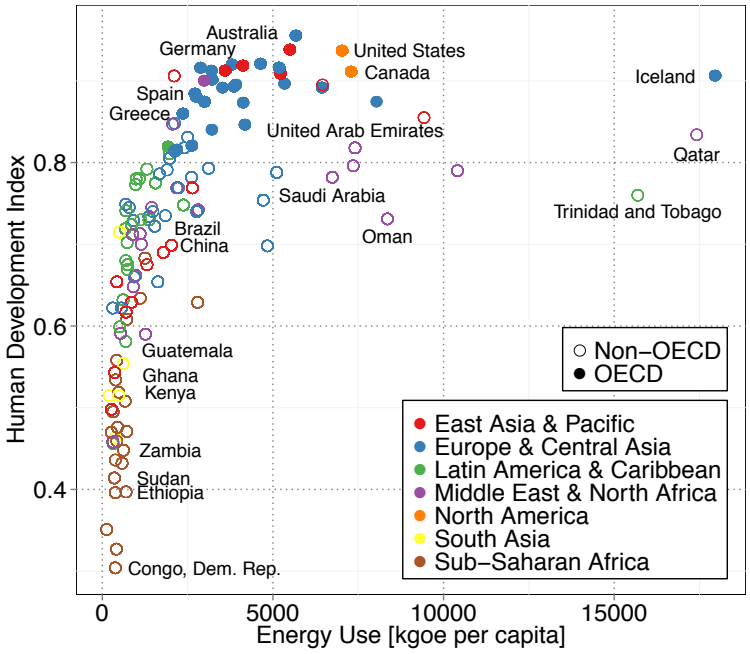

Figure 1.1: Human Development Index vs. energy consumption.

An interesting correlation exists between the HDI and per capita energy consumption as shown in Figure 1.1. Per capita energy consumption was gathered from publicly available World Bank data, and the HDI is published by the UNDP [98, 100]. A distinction has also been made between member countries of the Organisation for Economic Cooperation and Development (OECD) since a separate correlation appears appropriate. The relationship between per capita energy consumption and HDI has been described using logarithmic, third-order regression, and saturation curves by previous researchers, and shows that higher HDI's (i.e. more developed, higher quality of life) are often associated with higher levels of energy consumption [29, 70, 82]. The shape suggests that an upper limit may exist where consuming more energy does little to further development (i.e. the saturation effect), and that small changes in energy consumption may effect drastic change in the progress of developing nations [70].

Presently, a significant disparity exists among the development level of nations, and much of the world resides at lower levels of human and energy development. The pursuit of progress among nations seems inevitable, and with it will come greater demands for energy. The U.S.

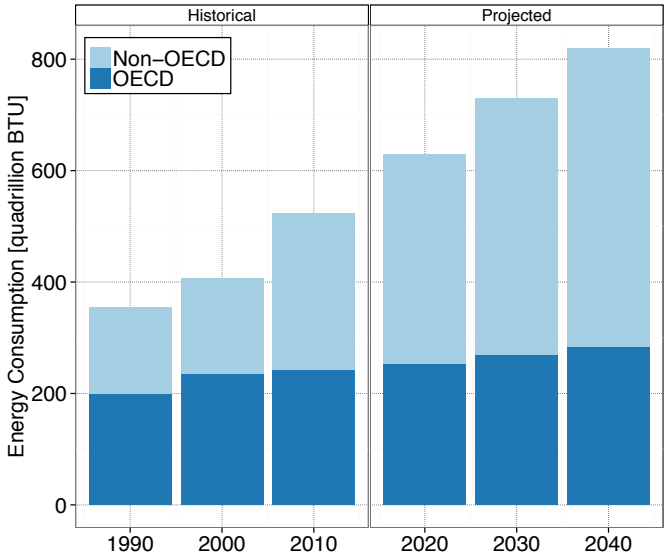

Figure 1.2: Projected growth in world energy consumption.

Energy Information Administration (EIA) currently projects that world energy consumption will increase to over 725 quadrillion Btus by 2030, a 40% increase from 2010 consumption [23]. The majority of growth is expected to be attributable to emerging economies (i.e. non-OECD) as shown in Figure 1.2. For reference, the distinction between OECD and Non-OECD countries is mapped in Figure 1.3. Ongoing concerns of fossil fuel depletion, climate change, energy security, and degrading environmental health prompt the question of how such expanding consumption can be achieved sustainably.

A first step towards a more sustainable future may be energy conservation (i.e. deciding to consume less energy). Energy conservation efforts play a valuable role in reducing wasteful practices and promoting necessary stewardship of limited resources. However, given the disparity in development among nations and projected energy growth, it is difficult to imagine the problem being solved entirely through energy conservation.

When considering the relationship between energy and humanity it is also important to consider issues of causation. Although correlation is apparent, causal relationships seem somewhat case dependent and ambiguous [79]. Soytas and Sari report causality running from energy consumption to GDP in France, Germany, Japan, and Turkey supporting the hypothesis that energy consump-

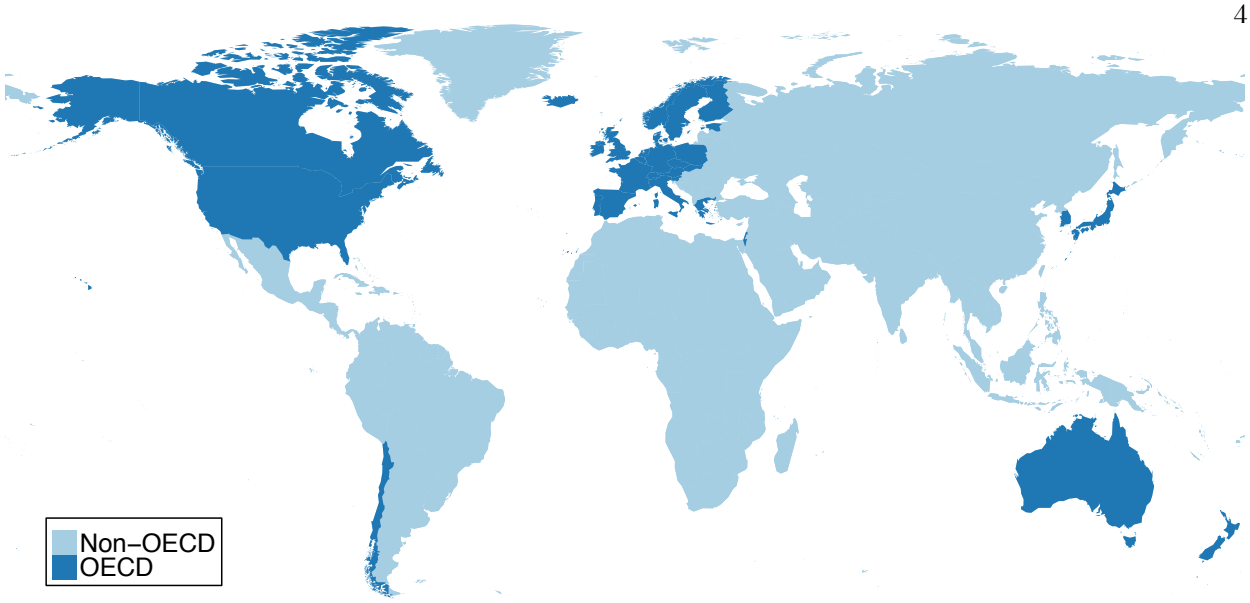

Figure 1.3: Map of OECD and Non-OECD countries.

tion fosters economic growth. Warr and Ayres also showed unidirectional causality running from energy consumption to GDP for the United States [103]. The reverse (i.e. causality running from GDP to energy consumption) was observed in Italy and Korea suggesting that increased GDP led to increased energy consumption. Bi-directional causality was found in Argentina, proposing an interrelated joint growth [96]. Additionally, Ouedraogo found negative cointegration between per capita energy consumption and HDI and positive cointegration between per capita electricity consumption and HDI, advocating that fuel source and quality may be equally as critical as energy quantity in the process of human development [79].

It is important to consider the nature of causality because it has important implications for energy policy formulation and may lead to differing viewpoints of energy consumption. On one hand, if causality runs from GDP/HDI to energy consumption, large-scale energy conservation policies can be implemented without fear of retarding progress. On the other hand, causality running from energy consumption to GDP/HDI would deter policies that promote curtailed energy consumption since stifled development could result. Additionally, policies that ultimately lead to increases in energy consumption may actually be encouraged to stimulate and maintain growth.

Since previous research has reported case-dependent causality in either direction, it appears that achieving a sustainable energy future requires a multi-faceted approach that extends beyond conservation.

A second avenue towards sustainability is to become more productive with each unit of energy consumed, i.e. increase efficiency. Improvements in energy efficiency imply either increased output for the same input, or decreased input for the same output. End-use energy efficiency is often touted as one of the World's greatest resources since small changes at the point of consumption can translate into greater impacts at the energy source. An illustrative example, published in *Scientific American*, shows that saving one unit of end-use energy can result in ten less units required at the energy source, after accounting for the effects of compounding losses along the distribution chain [68]. The International Energy Agency (IEA) has recently quantified the cumulative impact of energy efficiency investments made 1974–2010 in eleven IEA countries to be 60 quadrillion Btus of avoided energy in 2010 [49]. This is approximately 65% of the 2010 total final consumption of the eleven countries in the analysis.

Increases in efficiency can result from from a variety of pathways such as technological advances, more intelligent consumption patterns, and energy recovery and reuse. Achieving higher productivity for each unit of input often requires great ingenuity and calls for wide-ranging solutions across disciplines. Furthermore, gains in efficiency should not be commingled with energy conservation. Efficiency improvements imply that the same level of service can be maintained at reduced input, whereas conservation bears the connotation of sacrificial service reductions and privation.

Therefore, it seems that advances in energy efficiency can withstand the causality debate surrounding consumption and development previously mentioned. Considering the case of causality running from GDP/HDI to energy consumption, efficiency increases would shift the relationship so that the same state of development may be realized at lower consumption. In the case of causality running from energy consumption to GDP/HDI, efficiency improvements could result in increases in HDI/GDP for the same level of energy consumption.

Although there seems to be sound reason to pursue energy efficiency, it is not without coun-

terpoint. The main argument against energy efficiency is often rooted in Jevons' paradox, or "the rebound effect", which holds that efficiency improvements actually result in increased demand since the relative cost of the resource is reduced by the gain in efficiency [80]. Greening, Greene, and Difiglio surveyed a variety of literature regarding this topic and conclude that while the rebound effect is not insignificant, the increase in consumption is generally much less than the savings created (i.e. a net benefit still exists from efficiency improvements) [37]. Similarly, a review by Sorrell, Dimitropoulos, and Sommerville concludes that the long-run direct rebound effect is likely less than 30% for efficiency gains in automotive transport, household heating, and household cooling in OECD countries [91].

Despite the potential for some reduction in savings due to "the rebound effect," it seems that energy efficiency is a worthwhile pursuit and can play a vital role in sustainably satisfying future energy demands.

1.2 The Opportunity for Efficiency in Buildings

Worldwide, buildings account for over 100 quadrillion Btus (32%) of final delivered energy and consume more energy than any other sector [88]. The EIA also projects building energy demand to increase faster than any other sector at an annual average rate of 1.6% [23]. As with global energy consumption, the majority of growth is expected in developing nations as economies expand and transition to modern facilities. For these reasons, the buildings sector is a prime candidate for effecting significant global impacts through advances in energy efficiency.

The opportunities for efficiency gains and improvements in buildings are near endless since buildings are a complex assembly of interacting passive and active subsystems. At the component level efficiency gains can be achieved through advances in technology that are inherently more productive in and of themselves. Examples of such advances can be observed in the development of light-emitting diode (LED) lighting technology and high efficiency refrigeration equipment. At the HVAC system level, gains can be achieved through the selection of appropriate system types (e.g. hydronic vs. air, variable vs. constant flow) and inclusion of heat recovery components. At

the structural level, façades, massing, and internal layout can be designed to appropriately leverage natural light, ventilation, and diurnal temperature variations. Thoughtful and purposeful integration of a variety of measures can result in significant improvements in whole-building performance. A technical report by the National Renewable Energy Laboratory (NREL) has suggested that large commercial office buildings can achieve 50% net site energy savings over minimally compliant (ASHRAE 90.1-2004) designs when efficiency measures are integrated holistically [61].

As building technology grows in complexity the harmonious orchestration of subsystems becomes paramount. Increasingly, intelligent control systems are required to realize the full potential of new technologies, e.g. daylighting designs should be accompanied by automatic lighting controls to avoid unnecessary use of electric lighting. Pacific Northwest National Laboratory (PNNL) has estimated that approximately 1 quadrillion Btus of primary energy can be saved annually through improved control of U.S. commercial buildings—nearly 6% of the total 2002 U.S. commercial building primary energy consumption [12].

1.3 Transitioning Energy Demands

Despite the significant gains that can be achieved through advances in energy efficiency, more will be required to sustainably address future world energy needs. Along with the increasing trend in overall energy consumption, the proportion of demand attributable to higher quality fuel sources is also growing. As previously noted, negative cointegration has been observed between energy consumption and the HDI [79]. In the subset of countries analyzed, 75% the total energy consumption was attributable to biomass (e.g. wood, wastes, dung), and negative cointegration suggests that consuming greater amounts of lower quality fuels can reduce developmental progress—presumably due to degradations in environmental health. The progression towards higher levels of development necessarily involves a transition to modern fuels and energy services.

The IEA projects that worldwide demand for electricity will grow more than any other final form of energy from 2011 to 2035. The expected growth rates range from 1.7% to 2.5% per year depending on energy policies implemented, and nearly 83% of the increased demand for electricity is

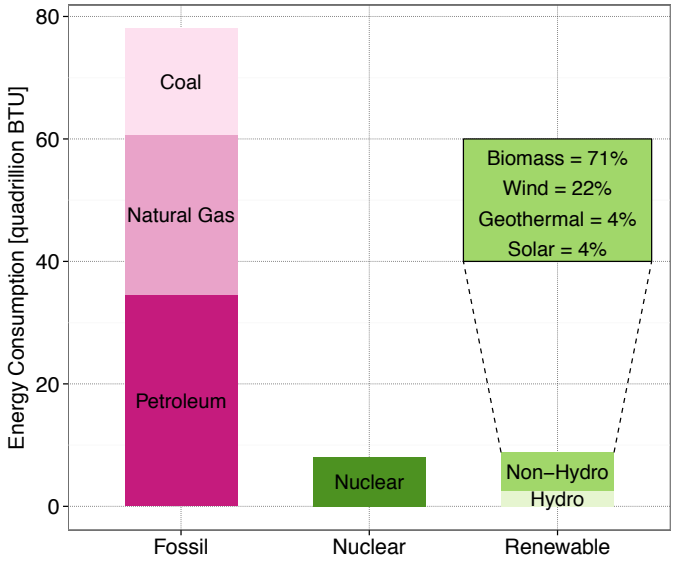

Figure 1.4: U.S. 2012 primary energy consumption by fuel source.

projected to be attributable to Non-OECD countries moving towards higher levels of development and transitioning to modern energy systems [9]. Although expected growth is smaller among OECD countries, developed nations are expected to continue their own transition toward greater amounts of renewable and lower carbon fuel sources (e.g. natural gas) [9]. In the U.S., 78 quads (82%) of primary energy consumed in 2012 was from fossil fuels, with the remaining 17 quads (18%) split between nuclear and renewables as shown in Figure 1.4 [22]. To create a more sustainable and secure energy future it seems a transition to greater renewables is necessary.

This transition, however, does not come without unique challenges. The U.S. electric grid, and many others around the world, has primarily been developed considering generation sources that behave quite predictably. Coal, gas, and nuclear generating plants all produce consistent output for known quantities of input fuel. A similar argument could also be made for hydroelectric power plants if any seasonal variations are well known. This consistency and predictability has historically allowed reliable scheduling of supply resources to accommodate changing and uncertain consumer demand. Predictable generation sources can be controlled and dispatched to maintain power system balance, ensuring a reliable supply of electricity.

In the U.S., the majority of prime hydroelectric resources are currently being utilized. Considering present technologies and economics, near future advances towards renewables are likely to be dominated by wind and solar resources. However, unlike traditional resources, the output of wind and solar generation is inherently variable making scheduling and dispatch of such supply sources seemingly impossible [95]. If not properly integrated, large quantities of variable generation may result in power quality issues, power flow imbalances, and grid stability issues [36].

1.4 Building-to-Grid Integration

One evident solution to accommodate the variable nature of wind and solar resources is to provide storage capacity. Grid storage technologies such as pumped-hydro, compressed air, and grid-scale batteries allow variable generation to be captured when the resource is available and dispatched as demand necessitates. Additionally, a NREL report studying the integration of wind and solar resources in the western U.S. determined that pursuing demand response and load participation programs would be critical in achieving higher penetrations of variable resources [35].

Historically, demand resources have played a relatively inflexible role in energy markets, requiring grid balancing operations to be achieved solely through modulation of generating resources. However, demand response can create additional grid flexibility to aid in absorbing the intermittency of variable generation resources. Many buildings contain significant amounts of thermal capacitance inherent in their construction materials that can be used to store thermal energy and buffer intermittent or varied HVAC operation. This thermal storage can also allow buildings to alter their energy consumption patterns to better align with supply resources without compromising indoor comfort. Considering the thermal storage potential of the existing U.S. building stock reveals a vast existing distributed storage resource that remains largely untapped.

Similar to traditional storage technologies (e.g. pumped-hydro), thermal mass can be charged and discharged considering various grid pricing signals, creating a real-time link between buildings and the electric grid. Real-time pricing signals can effectively communicate impending grid constraints and congestion, giving buildings an opportunity to provide relief if possible. Active

participation of buildings in energy markets is one way to directly integrate building and electric grid operations, achieving greater system level efficiencies through improved matching of supply and demand.

In addition, buildings may also be able to supply valuable ancillary services to the electric grid (e.g. frequency regulation and spinning reserves), alleviating the need to provide these services inefficiently using fossil fuel generation. Response from flexible loads may be able to provide ancillary services in a more accurate, reliable, and prompt manner than traditional generation equipment [55]. Including faster resources in the mix may ultimately reduce the total amount of ancillary services required without compromising reliability. The grid benefits from lower operating costs while loads generate revenue through providing grid services. Additional benefits of incorporating flexible load response into energy markets may also include: increased system reliability, improved market efficiency, risk management, reduced environmental emissions, market power mitigation, and increased system efficiencies [52].

In general, previous work has evaluated the potential for commercial buildings to participate in each energy and ancillary service market separately. Ideally building operation should be planned in consideration of all viable markets to maximize the benefits to both building owners and the electric grid. An important principle of maintaining separate markets for energy and individual ancillary services is that price signals can better communicate the needs of the grid to suppliers and consumers. Optimizing building operation in the context of multiple markets creates an opportunity for buildings to respond to grid needs when feasible and economically attractive. Including both ancillary service revenues and energy cost within the optimization allows for determination of an operational strategy that balances the severity of grid needs with the desire for lower utility bills (or other objectives such as reduced energy consumption).

1.5 A Portfolio Perspective

Traditionally, commercial buildings have been viewed as independent entities consuming grid resources to meet the needs of their respective tenants. Past work in optimal control of buildings

has considered a single building's ability to shift load based on hourly real-time prices and demand charges, resulting in facility level utility savings. This operation, while optimal in the individual sense, neglects the fact that buildings are all connected to the same electric grid. The aggregation of individual optimal solutions may in fact be suboptimal when considering the characteristics and operations of other buildings and viewing the problem from a communal perspective.

Buildings are diverse in physical design and operation and it seems a higher-level of coordination may be necessary to achieve the full benefit of building-to-grid integration by unlocking opportunities for cooperation on achieving joint objectives. More than simply providing a demand response mechanism, it is suggested in this research that model predictive control of building portfolios provides a framework for optimally managing multi-building load resources such that greater benefits can be provided to building owners and the electric grid than when optimizing buildings independently. By giving the optimizer the knowledge of all unique building characteristics available within a portfolio of buildings, various features may be exploited to orchestrate an optimal combined operation of all portfolio members.

Fundamental to the idea of this research is the belief that diversity among building characteristics and operations creates an opportunity for synergy. Building system interactions are complex and often difficult to comprehend for a single building, which makes it difficult to know how and when synergistic effects may arise among building portfolios. Therefore, further research is needed to investigate the opportunity for synergistic effect among building portfolios and motivate future pursuits of cooperative load control.

1.6 Research Questions and Objectives

Overall, this research sought to explore the opportunities for, and benefits of, multiple buildings simultaneously participating in multiple electric grid markets. Several specific questions were of interest regarding this subject:

(1) Participation in ancillary service markets requires estimating the potential capability of a resource and bidding availability into markets in advance of the operating hour for optimal dispatch of resources. How can a commercial building's ancillary service capacity be estimated and its availability to participate be forecast?

(2) Providing ancillary services has the potential to generate revenue for buildings, while load shifting and real-time price response creates utility cost savings. Can building operations be optimized to simultaneously participate in multiple ISO markets? Are there mutual benefits to buildings and electric grid?

(3) Previous work has shown the potential benefits of load diversity within a portfolio. Are there opportunities for synergy through model predictive control of building portfolios? If so, when and how do they materialize? What level and forms of diversity must exist within the portfolio for benefits to be achieved?

(4) Are certain buildings better at certain tasks? Can portfolios provide services unattainable by single buildings?

(5) The amount of work performed by a portfolio may not be distributed evenly among its members. Is aggregation still beneficial if it causes certain buildings to perform more work?

These questions were investigated through the following primary objectives:

(1) Develop reduced order models to enable multi-building optimization. Demonstrate the performance of reduced order models for supervisory model predictive control of building thermal mass and systems.

(2) Develop a methodology for estimating building ancillary service capabilities.

(3) Extend MPC environment to multi-market optimization (e.g. simultaneously optimize participation in real-time energy and frequency regulation markets).

(4) Extend MPC environment to multi-building optimization and explore ensemble synergy through simulation. Juxtapose results with individual optimizations.

(5) Explore multi-market, multi-building, optimization in consideration of non-optimized buildings.

1.7 Organization

This dissertation has been organized into the following chapters:

- Literature Review: This chapter provides a survey of literature surrounding topics related to this research.
- Methodology: This chapter describes the high-level approach to exploring opportunities for synergistic effect among building portfolios.
- Reduced-Order Modeling and Simulation Environment: This chapter discusses the development of a reduced-order modeling environment necessary to pursue portfolio simulation studies.
- Multi-Market Optimization: This chapter describes the methodology developed to estimate the frequency regulation capability of a commercial office building, and presents simulation studies co-optimizing building participation in multiple grid markets.
- Portfolio Optimization Development and Testing: This chapter describes the extension of the building optimization to multiple buildings, as well as presents some initial portfolio optimization results.
- Portfolio Case 1: Retail and Large Office Buildings: This chapter provides an in-depth discussion of a portfolio simulation case considering 120 retail buildings and one large office building.

- Portfolio Case 2: Medium and Large Office Buildings: This chapter presents an in-depth discussion of portfolio optimization results for a portfolio of 20 medium office buildings and one large office building.

- Portfolio Case 3: Medium and Large Office Buildings with High Demand Limit:
  This chapter provides a third portfolio case study that places greater emphasis on frequency regulation within the portfolio of 20 medium office buildings and one large office.

- Conclusions and Discussion: This chapter summarizes and discusses the primary findings of this research and provides some insight for future work.

Chapter 2

Literature Review

2.1 Building Energy Modeling for Control Applications

A wide-variety of building modeling techniques have been developed throughout history to serve the various needs to quantify building energy performance. The strategies range from steady-state calculations to detailed dynamic computer simulations. Various factors such as accuracy, sensitivity, versatility, speed, cost, reproducibility, and ease of use should be consider in selecting a modeling approach. In the context of building control applications, a method that produces relatively high resolution estimates (i.e. hourly or sub-hourly) and high accuracy is often desired. Simple steady-state models such as degree-day, change-point, or bin methods are typically unable to capture the necessary dynamics for control applications and are recommended for longer-term diagnostics and energy savings calculations. For control, dynamic modeling methods such as artificial neural networks (ANN), thermal networks, and detailed computer simulation are recommended. Dynamic modeling methods may be further classified as either being "forward" (white-box) or "data-driven" (black-box). Forward methods use known physical parameters and relationships to compute outputs from inputs, while data-driven methods seek to identify relationships using statistical data analysis [4].

2.1.1 Forward Modeling

Computational advancements have led to the development of numerous forward modeling software tools that vary in form and functionality. Crawley et al. provide a detailed comparison of

twenty major simulation programs in [25] with a summary also provided in [26]. Factors such as zone load calculations, convection algorithms, daylighting and solar models, and HVAC modeling capabilities are contrasted. A few of the more popular in use in the U.S. today are TRNSYS, DOE-2, and EnergyPlus.

The TRNSYS simulation engine solves the system of algebraic and differential equations that is constructed by connecting pre-defined simulation components. The simultaneous solution of building and HVAC systems in TRNSYS has made it a good candidate for a variety of simulation tasks. Henze et al. have successfully use TRNSYS models to investigate building thermal mass control strategies in [32, 38, 44]. TRNSYS has also been used to validate other modeling approaches and provided surrogate data for inverse methods as in [1, 14, 15, 16]

DOE-2 is based on sequential simulation of building thermal loads and HVAC systems, with the load calculations making use of response-factor and weighting-factor methods [4, 60]. Although fairly quick from a computational perspective, this sequential, "single-pass," format somewhat limits its ability to model advanced control strategies since there is no feedback between zone and system simulations.

The development of EnergyPlus further improved upon DOE-2 capabilities by implementing Heat Balance Method load calculations as well as integrating the simulation of building elements and HVAC systems through an iterative solution scheme. The integrated simulation approach resulted in more realistic modeling of system controls, system capacity limitations, moisture adsorption and desorption in building elements, radiant conditioning systems, and space temperature predictions [97]. EnergyPlus load calculations also abandoned the response factors in favor of the conduction transfer function formulation [31, 89]. EnergyPlus has been successfully used to evaluate advanced building control methods including optimal control of mixed-mode building operation in [72] and optimal HVAC demand response strategies in [75]. It is worth noting that these studies focused on optimal supervisory control strategies of hourly time discretization. While EnergyPlus can simulate at sub-hourly time steps, the component models are often based on quasi-steady state formulations that have limited ability to model short term transient effects.

Of more recent development in regard to building simulation is the Building System Library being developed for Modelica. Modelica is an equation-based object-oriented modeling language that has grown acceptance primarily among the automotive industry. The equation based format allows components to be described using discrete, algebraic, and differential equations for correct modeling of transient effects. Components are also constructed to represent a physical device where models interface as they would in reality. This characteristic leads to straight forward system modeling based on its physical manifestation as well as object and code reusability. The system of equations generated by model connections can be symbolically manipulated and paired to the most appropriate numerical solvers for the particular task. Decoupling the solvers from simulation routines can reduce numerical noise which can be an important factor when coupling simulation with optimization [104, 105, 106].

2.1.2 Inverse Modeling

At the other end of the spectrum, a wide range of inverse modeling techniques have also been applied to energy modeling of buildings. Wu, Reddy, and Claridge describe traditional multiple regression techniques and principal component analysis in [107]. Dhar, Reddy, and Claridge also applied Fourier series modeling to predict hourly energy use of buildings noting that building load and consumption profiles are often strongly periodic. Good agreement was observed when data were categorized by day-type to remove operational differences and separate models were used to predict weather dependent (e.g. heating and cooling) and independent (e.g. lighting) loads. Concluding remarks recommend this method for retrofit savings analysis, diagnostics, and potentially short-term forecasting [28].

Krarti et al. have shown neural networks are able to predict building retrofit savings at hourly resolution in [57]. Trained nets did not extrapolate well to other climate zones and it was recommended that training be performed for each specific building. Krarti also applies neural networks to weather forecasting, next-day load predictions, ground heat transfer modeling, and control of an active thermal energy storage system in [56]. Good performance is shown for the

applications, and it is noted that further research may be necessary to improve reliability, accuracy, and computational efficiency before wide spread implementation is feasible. Successful application of neural networks to building energy prediction are also demonstrated by Dodier and Henze in [30], and Karatasou et al. in [51].

The inverse methods discussed, thus far, can also be described as pure "black-box" models in the sense that the mapping from input to output is very complex with little or no information pertaining to building physics. Various hybrid approaches also exists, coined "gray-box", that blend the benefits of inverse and forward modeling techniques. Inverse gray-box methods typically rely on a model formulation that is rooted in physical relationships with certain physical parameters being identified inversely from measured data.

As with several forward modeling simulation tools the roots of this work extend back a few decades to developments in computational heat transfer. Ceylan and Myers showed heat transfer functions could be derived from a set of first order differential equations [18]. Seem et al. then demonstrated the combination of building element transfer functions into a comprehensive room transfer function (CRTF) to model heat transfer to a zone [89]. The convenient form of heat transfer functions relates outputs at the current time step to a sum of weighted current inputs, and past inputs and outputs. It was a natural extension to apply regression techniques to estimate model coefficients rather than directly computing from known characteristics. In the early 90's, Braun presented a simplified version of Seem's comprehensive room transfer function where equation coefficients were estimated through nonlinear regression between transfer function predictions and data from a detailed TRNSYS model. Transfer functions were then used to predict zone sensible loads and temperatures in the investigation of optimal control of building thermal mass [14].

This method was further improved upon in later work by Chaturvedi and Braun who extend the inverse parameter identification to a zone thermal RC network [15, 17, 19]. Since thermal resistance/capacitance (RC) networks can be reduced to a system of linear, time-invariant, differential equations the approach described by Seem [89] can be used to convert to a transfer function representation. The inverse parameter identification is improved by placing physical constraints on

R and C parameters. In effect, measured data is used to estimate physically plausible parameters of a physically structured model. Lee and Braun use similar inverse gray-box model formulations to investigate demand limiting thermal mass control strategies in [62, 63, 64, 65, 66].

Thermal RC networks also benefit from the fact that they need not be inverse models. Since the models are based on an electrical analogy of heat transfer physics RC parameters may be assigned based on material specific heat transfer coefficients. Often total resistances may be known wall systems and frequency domain regression is performed to optimize nodal placement as in [77, 101, 102, 114].

Armstrong has also applied regression techniques to identify transfer function coefficients in [2]. Model identification is performed with overall U-Factor constraints as well as pole constraints in the frequency domain. The purpose of these additions was to improve the reliability and stability of parameter estimates. Armstrong also makes use of a total least squares objective function to minimize both load and temperature predictions. Application of the models is addressed in a companion publication [3].

2.2 Thermal Mass Control

Thermal mass control of buildings has become a popular topic within the last few decades. In the early 90's, Braun showed optimal thermal mass control had the potential for significant savings by applying dynamic optimization to computer simulations in [14]. Rabl and Norford investigated night pre-conditioning to reduce peak demand and describe concepts such as "storage efficiency" and "COP variation" in [85].

Morris et al. used simulations to develop minimum energy cost and minimum demand pre-cooling strategies. Strategies determined through simulation were applied to a test facility and about 10% energy savings and 38% demand savings were observed [73]. A similar precooling strategy was developed for a 1.4 million square foot office building by Keeney and Braun in [53]. In this case, a strategy was tested to limit the peak cooling load for continued building operation in the event of the loss of one of four 900 ton chillers. Field tests showed peak load was successfully

limited to 75% when an appropriate precooling strategy was deployed.

Braun continued to investigate thermal mass control topics into the 2000's. Braun and Montgomery used simulation to investigate the performance of several precooling strategies including: light precool, moderate precool, extended precool, slow linear discharge, fast linear discharge, and maximum discharge in [17]. No optimization was used in these studies however precooling strategies were tested for a variety of climates and utility rate structures. Potential savings were significant however it was also noted that savings were dependent on climate, utility rate structure, and building properties. Braun and Lee used parametric simulation to evaluate demand limiting (DL) strategies in [62]. Demand limiting strategies incorporated pre-occupancy precooling and maintained the lower occupancy temperature bound until the start of the peak demand period. Similar DL strategies were successful in field experiments by Xu et al. in [108]. Braun and Lee further developed thermal mass DL strategies by proposing practical methods for determining building specific DL setpoint trajectories in [63, 64, 65]. The work was then further extended to include model-based optimization for determining DL setpoint trajectories in [66]. It should also be mentioned that the basis for much of this work utilized the inverse gray-box modeling techniques developed previously in [15, 19].

Henze has also contributed significantly to thermal mass studies with a string of publications throughout the 2000's. Henze et al. investigated simultaneous optimal control of active and passive thermal storage in [45]. A two stage optimization was used to first determine zone setpoints that optimized the daily cooling load (passive storage), and second to optimize use of active storage to further minimize the cost function. The observed combined savings was greater than either technology alone, however, was not greater than the sum of independent savings. Henze also explored optimizing energy cost compared to energy savings. It was determined that passive and active storage were important factors in reducing overall cost, however neither were used when minimizing site energy consumption [43]. The active and passive optimal control method was field tested in a real-time experiment in [47]. The experiment implemented the model-based optimization technique with BAS communication procedures to provide building control over a five

day period. Savings were observed despite imperfect weather predictions and model mismatches, however achieved savings were lower than expected. Further investigation concluded that the test facility may not have been optimal for passive thermal storage and that realistic plant models may be necessary to account for part-load characteristics. In an effort to characterize ideal candidates for passive thermal mass strategies, Henze et al. performed an extensive investigation to assess the sensitivity of optimal mass control with respect to utility rate structure, occupancy, internal gains, building mass, occupied temperature setpoint range, and weather. Key findings are summarized in [48]. Efforts were also made to determine near-optimal guidelines based on various factors in [46]. Henze et al. also describe a methodology for the inclusion of demand charges in the optimization task through the addition of a penalty function that responds to exceedance of a"target demand limit" [44]. Greensfelder and Henze also demonstrate the extension to real-time pricing in [38].

2.3 Frequency Regulation Ancillary Service

Operating a reliable and effective large-scale electric power system requires the procurement and scheduling of resources over several time scales. Long-term planning secures the availability of adequate generation capacity to meet changing consumer needs, while short-term scheduling and dispatch activities ensure real-time stability through continuously balancing electricity supply and demand. With respect to short-term operations, modern energy markets often implement a forward (e.g. day-ahead) market to schedule adequate resources for the next operating day, and a real-time market to balance actual intra-hour conditions experienced during operation. In addition to the scheduling of generation capacity, various ancillary services, such as frequency regulation, spinning reserves, voltage control, and black start, are required to maintain power system reliability [42]. Frequency regulation ancillary service is responsible for correcting small deviations between supply and demand that occur on minute or sub-minute time scales, while spinning and non-spinning reserve services provide a safeguard against generator failures and contingency events. Determining the optimal assignment of energy supply and ancillary service duties among competing providers is nontrivial, and has typically been approached by seeking the lowest operating cost solution that

maintains security and reliability requirements [78]. In order to be considered in economic unit commitment and dispatch, resources must determine their cost of providing various grid services and submit quantity and price bids to the independent system operator (ISO).

2.3.1 Traditional Generation

The concepts surrounding optimal bidding and scheduling have been well developed for traditional generation and storage resources [81]. Dynamic programming approaches have been proposed for scheduling large pumped hydro storage in coordination with thermal unit generation [20]. Day-ahead hydro and thermal generator bids have been determined using a mixed-integer linear programming formulation with stochastic day-ahead price forecasts and risk constraints [34]. Unit commitment of pumped hydro considering the effect of reservoir head on power production is presented in [10], and the scenario of scheduling multiple, independently-owned, hydropower plants operating from the same stream flow is considered in [86].

2.3.2 Vehicle-to-Grid

Recent work towards grid integration of electric vehicles also provides an important source of relevant literature. While attached to the electric grid, electric vehicle batteries can provide flexible electric storage for use in providing ancillary services. Intelligent charging algorithms for providing frequency regulation have been developed [40, 93], as well as a methodology for scheduling both spinning reserves and frequency regulation [92]. An extension was also made to consider selling energy back to the grid (i.e. battery discharging), unexpected vehicle departures, and battery degradation costs due to cycling [94]. In general, formulations seek to determine a preferred operating point (i.e. baseline charging or discharging strategy) and the additional power draw limits that maximize the profits of the electric vehicle aggregator.

2.3.3 Commercial Buildings

With respect to the building science domain, previous work has demonstrated that significant peak electric demand reductions can be achieved through utilizing building mass as a thermal storage medium [14, 17, 109]. Building energy simulation programs have been coupled with optimization routines in a MPC framework to determine zone temperature setpoint strategies that minimize building utility cost considering time-of-use electric rates and peak demand charges [44]. Real-time pricing scenarios have also been considered [24, 38], creating a link between building operations and energy market pricing signals. The traditional temperature-based MPC problem has also been expanded to include building-wide interactions among CO_2, humidity, and pressure to exploit additional degrees of freedom and storage mediums [110].

Aside from the benefits achieved through load shifting and price response, buildings with significant thermal mass may also be well-suited to provide ancillary services since zone thermal inertia can buffer intermittent or varied HVAC operation [42]. Residential air conditioners have been evaluated for providing spinning reserve services [58], and a pilot study successfully bid commercial building demand response into day-ahead non-spinning reserve markets [54].

Recent work has also considered controllable building electric loads for economic dispatch in energy markets. As an example, chilled water supply temperature was modified to create changes in electric demand [8]. The inverse chiller model used computes chiller power as a function of chilled water supply temperature [7], and the approach made use of model linearization and simplifying assumptions to formulate a linear programming problem. Zone thermal dynamics were not explicitly modeled and chilled water temperature setpoints were discretized in coarse six hour blocks. The work was extended to also include transmission constraints [6].

Frequency regulation (FR) in commercial buildings has also recently been investigated [111]. An overview of the FERC Order 755 "pay-for-performance" rule, as implemented in PJM, was provided highlighting the opportunity for buildings to be compensated for faster and more accurate response to regulation dispatch. Detailed thermal zone and HVAC models were used to evaluate

building response while tracking a FR signal. In one case, the FR signal was used to modulate the duct static pressure setpoint of the building air handling unit, achieving a primary power response due to changes in fan speed. In a second case, FR was achieved through modulating zone temperature setpoints, achieving whole HVAC system response due to changes in zone cooling load. Related work has also proposed combining the regulation signal with a variable speed fan control signal to directly modulate fan speed [41].

2.4 Load Aggregation

ASHRAE RP-1146, titled "Building Operation and Dynamics Within an Aggregated Load" sought to identify situations when it is beneficial to manage total multi-building demand through load aggregation and investigate control strategies that would further reduce energy costs at the aggregate level. Load aggregation seeks to benefit from the fact that coincident load peak is not necessarily the sum of individual load peaks (i.e. diversity). One aspect of this work included a simulation study of an office, retail, and hotel building to illustrate load aggregation benefits as well as explore the combined effects of load aggregation with curtailment measures. Load curtailment included simultaneous lighting power, equipment power, and ventilation rate reductions as well as temperature setpoint increases. Simulation results showed that: 1) approximately 8% of demand cost savings could be achieved through load aggregation, 2) 6% demand cost savings were possible when aggregating loads with curtailment measures already in place, and 3) 36% demand cost savings could be achieved through the combined addition of curtailment measures and load aggregation [87].

Chapter 3

Methodology

This chapter serves as an overview of the steps taken to complete the research objectives, and answer the questions presented in Chapter 1. A more detailed description of each method is provided in subsequent chapters.

Objective 1: Develop reduced order models.

Two methods of reduced-order modeling were explored: 1) International Standard ISO 13790 "Simple Hourly Method" and 2) inverse gray-box thermal RC networks. The inverse gray-box reduced order modeling approach was chosen due to its ability to accommodate flexible model structures, as well as utilize measured building data. Two methods were also explored for parameter estimation of inverse gray-box model parameters: 1) nonlinear least squares and 2) Bayesian parameter estimation. Although Bayesian methods are more powerful for a variety of reasons, it was determined that the objectives of this research do not fully exploit these advantages and traditional methods were utilized for the majority of model identification tasks. A model complexity analysis was performed considering six RC models to determine an appropriate thermal network structure for retail, medium office, and large office building models. Corresponding system models were developed for each building based on quasi-steady state component used in popular white-box building energy simulation programs (e.g. EnergyPlus). Three levels of model validation were performed corresponding to the component, system, and whole-building levels to ensure appropriate performance.

Objective 2: Develop methodology for estimating building ancillary service capabilities.

To estimate the capabilities of a commercial building to provide ancillary services a model perturbation approach was taken. A whole building energy model was used for its ability to capture important relationships between building and equipment operating constraints. Frequency regulation was assumed to be provided through zone setpoint or chilled water setpoint modulation. The whole building model was used to explore various setpoint changes and observe the power response that may be utilized for frequency regulation. For frequency regulation via chilled water setpoint modulation, two methods were developed that characterize different responses (i.e. fast response and steady response). Estimating the potential power change for frequency regulation allows the potential ancillary service revenue to be computed as well.

Objective 3: Extend MPC environment to multi-market optimization.

The MPC environment was extended to included multi-market optimization by embedding the frequency regulation estimation methodology within the controller. As part of each function evaluation, the frequency regulation potential is estimated using the perturbation approach so that regulation revenue can be included in the MPC objective function. This allows the optimizer to weigh setpoint strategies that allow for increased regulating potential with energy cost and demand limiting objectives. Simulation case studies were performed for the medium office and large office buildings under both high and low demand limits.

Objective 4: Extend MPC environment to multi-building, multi-market optimization.

The single building multi-market optimization was extended to accommodate multiple buildings. Overall, this was performed by generalizing subtasks within the controller to be performed for any number of buildings. The cost function is computed based on the aggregated results of

all buildings, producing a single cost for the portfolio-wide control solution. Buildings within the portfolio framework may be subject to different weather, energy pricing, optimization variables, discretization, and simulation engines.

Objective 5: Explore multi-market, multi-building, optimization including non-optimized buildings.

The multi-building optimization environment was used to perform 18 portfolio case studies with varying portfolio constructions and optimization parameters. Synergistic effect was identified by comparing the portfolio optimization results to the aggregated results of optimizing each building individually. Portfolio cases included scenarios where one building may be optimized in consideration of unoptimized buildings.

Chapter 4

Reduced-Order Modeling and Simulation Environment

4.1 ISO 13790 Building Model

As an initial candidate for reduced-order modeling an evaluation of ISO13790 was performed. International Standard ISO 13790 [50] provides a detailed methodology for calculating the energy use for space heating and cooling in buildings. Prescribed load calculations are intended for use in building design rating, energy performance comparison, retrofit impact analysis, and regional prediction of future energy resource demands. The standard was developed, among others, to orchestrate European energy performance calculations, and outlines three calculation methods of varying resolution and detail. The methodology under evaluation here follows the Simple Hourly Method. The underlying principles of the Simple Hourly Method are based on the five resistor one capacitor lumped parameter model shown in Figure 4.1. Resistances are defined for transmission through opaque materials and glazing ($H_{tr,em}$, $H_{tr,ms}$, and $H_{tr,w}$), convective exchange from interior surface to the zone air ($H_{tr,is}$), and direct exchange between ventilation/infiltration and the zone air node (H_{ve}). All internal gains due to occupants, equipment, and solar radiation are distributed in fixed proportion between the mass, surface, and air nodes, and a single capacitance is designated for all building mass. Space conditioning is assumed to be provided through heating and cooling energy applied at the air temperature node.

The resistance network is solved via a Crank-Nicholson numerical technique and the resulting expressions are provided in the standard. Expressions are evaluated at each time step assuming no heating or cooling is provided to determine the free-floating temperature of the zone. The free-

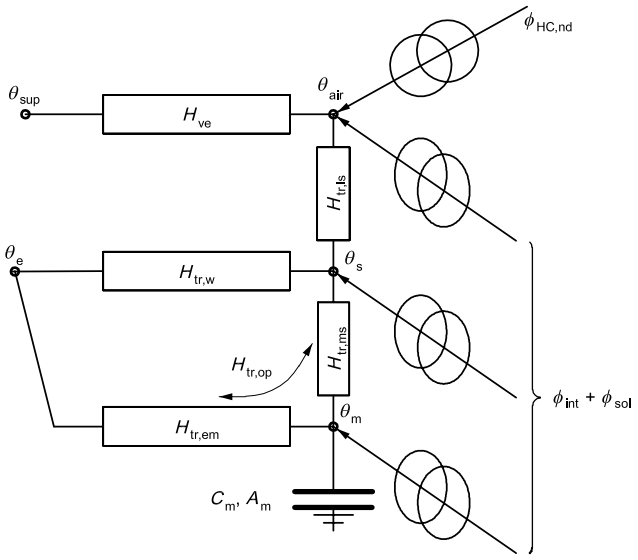

Figure 4.1: 5 resistor 1 capacitor model of ISO13790 Simple Hourly Method [50].

floating temperature is compared with heating and cooling set points to determine if heating or cooling is necessary. If heating or cooling is needed, a fixed heat input is assigned and a new zone temperature is calculated based on the heat input. The ratio of the temperature difference between free-floating conditions and set point and temperature change due to known heat input is equal to the ratio of actual heat needed to known heat input as shown by Equation (4.1).

$$\frac{q_{need}}{q_{assigned}} = \frac{T_{setpoint} - T_{float}}{T_{heatassigned} - T_{float}} \tag{4.1}$$

The methodology described by ISO 13790 was implemented in Matlab by Corbin (2010) to produce a computerized simulation environment. The initial implementation of the standard currently models a single zone, however the standard provides details for expanding the calculation to multiple thermally coupled zones as well.

For comparison and testing of the ISO model, the DOE Stand-alone Retail Commercial Benchmark Model was used as a test case [99]. Initially, the five zone EnergyPlus model was simplified to more closely align with ISO model characterisitics. Simplifications included: creation of single zone, substituting HVAC system with an ideal loads air system, and specifying radiant/convective fractions of internal loads to equal the standard. EnergyPlus model inputs were

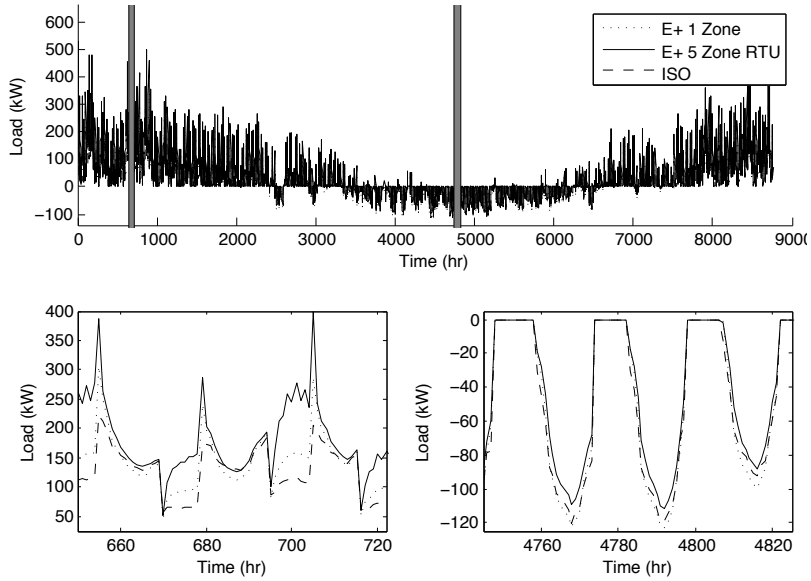

Figure 4.2: Load comparison of five zone EnergyPlus and ISO 13790 models

extracted from the IDF file and used in specifying parameters for the ISO model. After achieving agreement between single zone models, the single zone ISO model was additionally compared to an unaltered five zone DOE benchmark model.

Simulations using the respective modeling methods were performed using Chicago TMY3 weather data, and model performance was evaluated by direct comparison of annual load calculations. Figure 4.2 shows the load calculations for a selected heating and cooling period during the simulation. During the heating period, some visual discrepancy is noted. The ISO model appears to underestimate heating need during early morning hours and slightly overestimate loads during the afternoon hours. The model appears to perform better during the summer period when disturbances are less extreme and load requirements are mild. Similar performance is observed for milder winter periods as well.

Single zone model discrepancies are numerically summarized in Table 4.1. Reported cumulative percent errors show that the ISO model underpredicts load requirements during heating and cooling seasons as well as annually. The largest difference occurs during the heating season as expected. The mean bias error likewise reports an average underestimation of load requirements

ranging from 2 kW to 7 kW. With respect to the mean, variations were observed to range from 30 to 70 percent.

Similar results were also observed in comparing the ISO model with the full five zone retail model as shown in Figure 4.2 and Table 4.1. The winter period shows potentially significant variations, on the order of 100 kW, during early morning periods. Increased dynamics in the EnergyPlus model are also noted due to the inclusion of packed roof top HVAC units. Overall, the ISO version appears to underestimate heating need and over estimate cooling need compared to the E+ model. Visual assessment is confirmed by seasonal and annual cumulative percent errors and mean bias errors. Mean bias ranges from 4 kW to 22 kW, and CV-RMSE from 46% to 93% of the mean.

Overall, the ISO 13790 load calculation routine produced similar results when compared to EnergyPlus for the cooling season, which is of primary concern in this work. The benefits of the ISO model include significant reductions in simulation time as well as reduced modeling effort. The ISO simulations performed only required specification of daily schedules, construction U-values and surface areas, mass characteristics, and internal loads to produce results comparable to more detailed EnergyPlus simulations. The ISO model makes good use of minimal building information, provided that key parameters can accurately be estimated. The method used here, however, is still a completely forward method and calibration can still be sensitive to parameter values. This methodology is one option for a reduced-order modeling framework, and a second candidate is discussed in the following section.

4.2 Inverse Gray-box Building Model

As a second reduced order modeling approach, inverse gray-box models were explored. The inverse gray-box modeling developed for this research is largely based on methods described by Braun and Chaturvedi in [15, 19]. Inverse gray-box models are based on the approximation of heat transfer mechanisms by an analogous electrical lumped resistance-capacitance network. This approximation creates a flexible structure that allows the modeler to choose the appropriate level of

Table 4.1: Single zone summary metrics.

	Single Zone			Five Zone			
	Heating	Cooling	Annual	Heating	Cooling	Annual	Units
Cumulative % Error	-14.6	-6.9	-12.4	-34.3	22.4	-26.3	%
Mean Bias Error	-7360	-1680	-4340	-22400	4130	-16300	W
CV-RMSE	41.8	29.9	69.5	64.5	45.5	93	%

abstraction. Model complexity can range from representing entire systems with a few parameters, to modeling each heat transfer surface with numerous parameters. Depending on the model structure and complexity, model parameters can approximate physical characteristics of the system. The inverse approach allows resistance and capacitance parameters to be estimated from building operational data, somewhat simplifying the overall model generation and thermal calibration process.

Figure 4.3 shows the 21 parameter thermal network representation found to work well by Braun and Chaturvedi [15, 19]. Other forms have been developed and explored in this work and are described in Section 4.4. A separate 3R2C network is used to represent external wall, ceiling, ground, and internal wall heat transfer. Looking at the 3R2C network for exterior walls, for example, $R_{e,1}$ could be thought to represent a combined external convection and radiation coefficient, $R_{e,2}$ wall conduction resistance, and $R_{e,3}$ internal combined convection and radiation coefficient to the zone air node. Solar gains from opaque elements are represented by $\dot{Q}_{sol,e}$ applied to the external surface nodes (e.g. $T_{e,1}$ and $T_{c,1}$). Storage is neglected for glazing elements which are represented by a single resistance R_w. Solar gains directly entering the zone through glazing are distributed among internal partition nodes $T_{i,1}$ and $T_{i,2}$ as $\dot{Q}_{sol,w}$. Internal gains are split into convective and radiant fractions. Convective fractions are applied directly to the zone air node T_z as $\dot{Q}_{g,c}$. Radiant portions are applied to interior surface nodes $T_{e,2}$ and $T_{c,2}$ as $\dot{Q}_{g,r,e}$ and $\dot{Q}_{g,r,c}$, respectively.

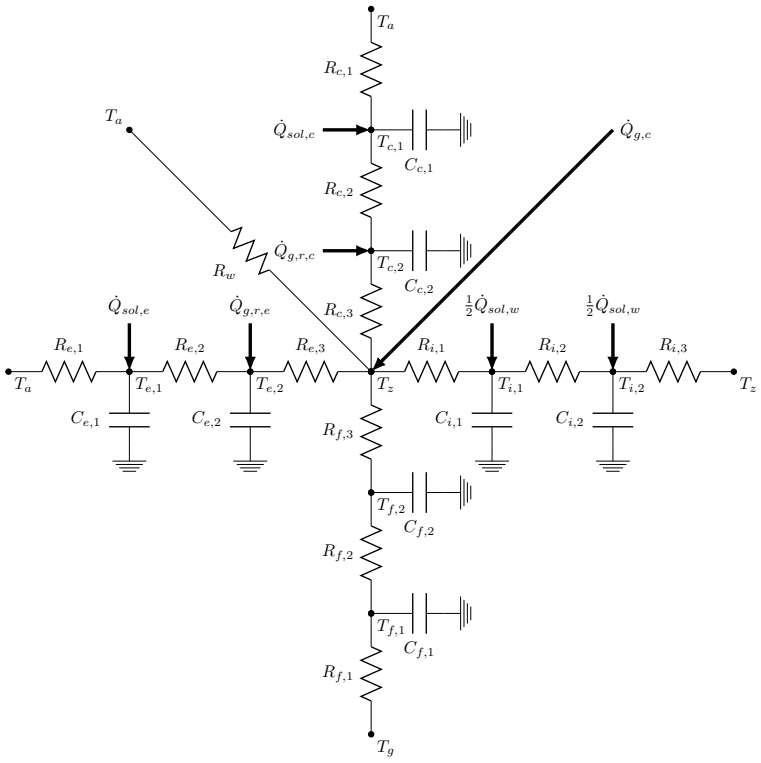

Figure 4.3: 21 parameter thermal RC network.

The thermal RC network can be represented by a system of linear first-order differential equations with constant coefficients by performing a energy balance at each node with a storage element. This system can be represented in traditional state-space form as

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$$

$$y = \mathbf{c}\mathbf{x} + \mathbf{d}\mathbf{u}$$

For the 21 parameter model the state and input vectors are

$$\mathbf{x}^T = [T_{c,1} \;\; T_{c,2} \;\; T_{e,1} \;\; T_{e,2} \;\; T_{f,1} \;\; T_{f,2} \;\; T_{i,1} \;\; T_{i,2}]$$

$$\mathbf{u}^T = [T_z \;\; T_a \;\; T_g \;\; \dot{Q}_{sol,c} \;\; \dot{Q}_{sol,e} \;\; \dot{Q}_{g,r,c} \;\; \dot{Q}_{g,r,e} \;\; \dot{Q}_{sol,w} \;\; \dot{Q}_{g,c}]$$

where T_z is the zone temperature setpoint, T_a is the ambient external temperature, $\dot{Q}_{sol,c}$ is the external solar gains incident on the roof, $\dot{Q}_{sol,e}$ is the solar radiation incident on exterior walls, $\dot{Q}_{g,r,c}$ is the radiant portion of internal gains applied to the ceiling surface node, $\dot{Q}_{g,r,e}$ is the radiant

portion of internal gains applied to the wall surface node, $\dot{Q}_{sol,w}$ is the solar radiation transmitted through glazing, and $\dot{Q}_{g,c}$ is the total convective internal gains. The state space equations are then converted to the following heat transfer function form

$$\dot{Q}_{sh,t} = \sum_{k=0}^{n} \mathbf{S}_k^T \mathbf{u}_{t-k\Delta\tau} - \sum_{k=1}^{m} e_k \dot{Q}_{sh,t-k\Delta\tau} \tag{4.2}$$

where $\mathbf{S}$ is a matrix containing input coefficients, e_k is a vector containing heat gain history coefficients, n is the number of past inputs in the calculation, and m is the number of past heat gain values in the calculation. The transfer function method is an efficient calculation routine as it relates the sensible heat gains to the space ($\dot{Q}_{sh}$) at time t to the inputs ($\mathbf{u}_t$) of n and heat gains ($\dot{Q}_{sh,t}$) of m previous timesteps. The input weighting coefficients ($\mathbf{S}_k^T$) and zone load coefficients (e_k) are the result of the state space to transfer function conversion process described by Seem [89].

Performing an energy balance on the zone air node in Equation 4.3 provides a basis for zone sensible load calculations where C_z is the zone air (or node) capacitance, T_z is the zone air temperature node, $\dot{Q}_{sh,t}$ is the zone sensible heat gain, $\dot{Q}_{inf}$ represents infiltration heat gain, and $\dot{Q}_{zs,t}$ is the required sensible zone load.

$$C_z \frac{dT_z}{dt} = \dot{Q}_{sh,t} + \dot{Q}_{zs,t} + \dot{Q}_{inf} \tag{4.3}$$

If the differential in Equation 4.3 is approximated by

$$\frac{dT_z}{dt} \approx \frac{T_{z,t} - T_{z,t-\Delta}}{\Delta}$$

it can be rearranged to develop an algebraic "inverse" transfer function for computing zone temperature predictions from a known zone load shown in Equation 4.4.

$$\bar{T_z} = \frac{\sum_{l=2}^{9} \mathbf{S}_0(l)\mathbf{u}_t(l) + \sum_{j=1}^{8} \mathbf{S}_j \mathbf{u}_{t-j\Delta\tau} - \sum_{j=1}^{8} e_k \dot{Q}_{sh,t-j\Delta\tau} + 2\frac{C_z}{\Delta\tau} T_{z,t-\Delta\tau} + \dot{m}_{inf} C_p \mathbf{u}_t(2) + \dot{m}_{SA} C_p T_{SA}}{2\frac{C_z}{\Delta\tau} - \mathbf{S}_0(1) + \dot{m}_{inf} C_p + \dot{m}_{SA} C_p} \tag{4.4}$$

In Equation 4.4, the $\dot{Q}_{zs,t}$ term from Equation 4.3 has been replaced with

$$\dot{Q}_{zs,t} = \dot{m}_{SA} C_p \left(T_{SA} - T_z\right)$$

so that zone temperatures can be computed based on the HVAC system supply air mass flow rate $\dot{m}_{SA}$ and supply air temperature T_{SA}. An assumption of this formulation is that the heat gains are computed using the average value over the time step so that the actual temperature at a given time step can be determined from

$$T_{z,t} = 2\bar{T}_{z,t} - T_{z,t-\Delta\tau}$$

Using the previous equations an ideal load calculation scheme for a dual-setpoint with dead band scenario can be described by the following procedure:

for t = simstart : simend do
  Calculate $\dot{Q}_{sh,t}$ using Equation 4.2.;
  Calculate $\dot{Q}_{zs,t}$ to maintain $T_z = T_{cool,set}$ using Equation 4.3 (assume cooling first).;
  if $\dot{Q}_{zs,t} < 0$ **(heating required to maintain cooling setpoint) then**
    Set $\dot{Q}_{zs,t} = 0$. Compute floating temperature using Equation 4.4.;
    if $T_z < T_{heat,set}$ **then**
      Re-compute $\dot{Q}_{zs,t}$ to maintain $T_z = T_{heat,set}$ using Equation 4.3
    end
  end
end

4.3 Parameter Estimation

Inverse gray-box models provide the opportunity to utilize measured building data to directly inform the model generation process. In this implementation models are trained using measured building sensible zone loads and temperatures to estimate the R and C values that best capture thermal dynamics. For air systems, sensible zone loads can be estimated from supply air temperature (or setpoint), supply air mass flow rate, and zone air temperatures via

$$\dot{Q}_{z,s} = \dot{m}_{sa} c_p (T_{sa} - T_z).$$

The parameter space can be constrained to consider only physically plausible values (e.g. no negative resistances), while allowing room for estimates to incorporate geometry and construction uncertainty as well (e.g. lower resistances could be chosen to compensate for an underestimation

of total external surface area). Table 4.2 shows the parameter bounds that are typically used in identifying the five parameter model.

Table 4.2: Reduced-order model training parameter bounds.

	R1	R2	R3	Rw	C1
Min	0	0	0.0303	0.05	78
Max	4.9886	4.9886	33.3336	3	536659
Units	m^2 K W^{-1}	m^2 K W^{-1}	m^2 K W^{-1}	m^2 K W^{-1}	J m^{-2} K^{-1}

The inverse gray-box implementation in this work also employs a zone capacitance multiplier on the zone air node to allow for characterization of mass objects that are in equilibrium with the zone air. Additionally an internal gains multiplier is used to scale the effects of internal loads on the zone air. When using surrogate data produced from known inputs this may be unnecessary, however, when utilizing real data there may be significant uncertainty regarding actual internal gains during the collection of training data. Including an internal gain multiplier allows the parameter estimation to tune the initial internal gain estimates if they are not well known.

4.3.1 Non-Linear Least Squares

Numerous researchers have considered the identification of thermal network models through parameter estimation using nonlinear least squares approaches or similar techniques [27, 15, 33, 2, 113]. It has been demonstrated that both time and frequency domain least squares minimization via traditional (e.g. Gauss-Newton) or metaheuristic (e.g. genetic) algorithms are beneficial. As a first approach to model identification, sum of squares minimization was used to identify model parameters that minimize the root-mean-squared error (RMSE), defined by Equation 4.5, between the reduced-order model predicted ($\dot{Q}_{rom}$) and surrogate or measured ($\dot{Q}_{ep}$) sensible zone load.

$$J = \sqrt{\frac{\sum_{i=1}^{N} (\dot{Q}_{rom,i} - \dot{Q}_{ep,i})^2}{N}} \tag{4.5}$$

In this work, a two-stage optimization was implemented that first performs a direct search over the parameter space to identify a starting point for local refinement. The direct search is performed on k uniformly random points located within the bounds of the parameter space. The local refinement, subject to the same parameter constraints, is performed via nonlinear least squares minimization implemented using the Matlab optimizer `lsqnonlin` based on trust-region Newton methods. The implementation in this environment also allows for local refinement to be performed around several good starting points from the direct search. It has been found that for higher complexity models the local optimization can be sensitive to the initial starting point. Good results have been found when the 12 best direct search points are given to separately executed least squares algorithms to simultaneously explore several attractive regions.

4.3.2 Bayesian Parameter Estimation

While least squares methods have been successfully used for inverse gray-box model identification, they typically result in point estimates of the parameters without consideration of uncertainties and can be sensitive to initial optimizer seeds or starting conditions. Lauret et al. has demonstrated improvements over traditional parameter estimation methods through the application of Bayes' Theorem to determine better estimates of convection coefficients for a radiant barrier roof system model [59].

As a second approach to identifying inverse gray-box model parameters required for this research, Bayesian inference was applied to explore the potential benefits of a probabilistic perspective. The Bayesian approach benefits over traditional methods because prior knowledge of the system is directly incorporated into the estimation task, and methods for addressing sensor noise are inherent to the Bayesian approach.

The probabilistic perspective maintains and utilizes all information available, and therefore can provide insight into the relationship between sets of model parameters, revealing tradeoffs and compensating interactions. The probabilistic approach also lends itself to continuous model uncertainty quantification and tuning where the posterior distribution of an initial parameter esti-

mation can be used as the prior for a subsequent parameter estimation update once new building performance data has been collected.

In this work, the Bayesian calibration of the gray-box model relied on the extension of a previously developed technique [74]. The inference can essentially be thought of as fitting a joint probability distribution to a measured data set. Specifically, conditional probabilities are related through the product rule to derive Bayes' Theorem and allow consideration of "before data" and "after data." The prior probability distribution (i.e. "before data") is updated with any measured data to form the posterior probability distribution (i.e. "after data"), which represents the state of knowledge in any inference task.

From a parameter identification perspective, the probability of parameter set Θ given measured data D and knowledge of the system K can be written as posterior probability $p(\Theta|DK)$. Bayes' Theorem then allows the conditional probability $p(\Theta|DK)$ to be computed from $p(\Theta|K)$, $p(D|\Theta K)$, and $p(D|K)$ as in Equation 4.6,

$$p(\Theta|DK) = p(\Theta|K)\frac{p(D|\Theta K)}{p(D|K)} \tag{4.6}$$

where $p(\Theta|K)$ represents prior knowledge about parameter values, $p(D|\Theta K)$ represents the likelihood of observing the measured dataset D given a particular parameter set Θ and knowledge of the system K, and $p(D|K)$ is the probability of observing the dataset. The relation can be written in alternate form where the numerator remains the product of likelihood and prior, and the denominator is a normalization factor so that posterior probabilities sum to unity.

$$p(\Theta|D) = \frac{p(\Theta)p(D|\Theta)}{\sum_i p(\Theta_i)p(D|\Theta_i)} \tag{4.7}$$

Assuming random Gaussian noise about a measured datum D_i, the likelihood of an observation can be determined from its location within the normal distribution with standard deviation σ_ϵ, centered at μ equal to the measured datum,

$$p(D_i|\Theta) = \frac{1}{\sigma_\epsilon\sqrt{2\pi}}\exp\left(\frac{-(D_i - M_i)^2}{2\sigma_\epsilon^2}\right), \tag{4.8}$$

where M_i is the model output given the parameter set Θ.

Assuming independent errors, the likelihood of the entire dataset is simply the product of the likelihoods of all individual points. The assumption is likely valid for common HVAC sensors (e.g. temperature probes), but correlated errors could be handled with a slightly different formulation—indicative of a fault model. Measurement errors are often correlated due to hysteresis error, linearity error, sensitivity error, zero shift error, and repeatability error. If such correlated errors are of concern, then a Bayesian (or other probabilistic) method may be used that can accommodate correlated measurements. The use of a least squares approach is more computationally efficient, it is however less robust with respect to many of the error sources found in sensor networks. This work, however, does assume uncorrelated temperature and sensible load measurements; thus it ignores autocorrelation of errors, which is estimated to be small. For the dynamics and time range considered in this problem, model structure is considered more important, with respect to data fit, than noise correlation. From this model assumption, we derive the easily computable likelihood function given by Equation 4.9, which happens to be equivalent to the least squares equation.

$$p(D|\Theta) = \frac{1}{\left(\sigma_\epsilon\sqrt{2\pi}\right)^n} \exp\left(\frac{-1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2\right) \tag{4.9}$$

Evaluating Equation 4.9 directly can pose numerical issues, since a small range of σ_ϵ values results in a large range of likelihoods. Double precision computing environments are typically capable of evaluating floating point numbers on the order of 10^{-308} to 10^{308}. This means that when using three weeks of hourly data (i.e. $n = 504$), σ_ϵ must approximately be in the range of $[0.1, 1.5]$. Values outside of this range will cause the likelihood (and consequently posterior) to evaluate to "Inf", "NaN", or "0" regardless of the time series fit. These numerical issues can be alleviated by computing the natural logarithm of the posterior rather than the posterior directly [59, 90].

To compute the natural log of the posterior, first, the log of both sides of the Equation 4.7 is

taken.

$$\ln\left(p(\Theta|D)\right) = \ln\left(\frac{p(\Theta)p(D|\Theta)}{\sum_i p(\Theta_i)p(D|\Theta_i)}\right) \tag{4.10}$$

The right hand side of Equation 4.10 can be separated using logarithm product and quotient rules.

$$\ln\left(p(\Theta|D)\right) = \ln\left(p(\Theta)\right) + \ln\left(p(D|\Theta)\right) - \ln\left(\sum_i p(\Theta_i)p(D|\Theta_i)\right) \tag{4.11}$$

The log-liklihood term of Equation 4.11,

$$\ln\left(p(D|\Theta)\right) = \ln\left(\frac{1}{\left(\sigma_\epsilon\sqrt{2\pi}\right)^n}\exp\left(\frac{-1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2\right)\right) \tag{4.12}$$

can be further simply by applying product and quotient rules as shown in Equations 4.13 and 4.14, respectively.

$$\ln\left(p(D|\Theta)\right) = \ln\left(\frac{1}{\left(\sigma_\epsilon\sqrt{2\pi}\right)^n}\right) + \ln\left(\exp\left(\frac{-1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2\right)\right) \tag{4.13}$$

$$\ln\left(p(D|\Theta)\right) = \ln\left(1\right) - \ln\left(\left(\sigma_\epsilon\sqrt{2\pi}\right)^n\right) + \ln\left(\exp\left(\frac{-1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2\right)\right) \tag{4.14}$$

It is known that $\ln(1) = 0$, and the power rule can be applied to the middle term of the right hand side. The last term of the right hand side simplifies due to logarithmic identity to produce Equation 4.15.

$$\ln\left(p(D|\Theta)\right) = -n\ln\left(\sigma_\epsilon\sqrt{2\pi}\right) + \frac{-1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2 \tag{4.15}$$

Recombining the simplified log-liklihood of Equation 4.15 with the log-posterior equation of Equation 4.11 produces:

$$\begin{aligned}\ln\left(p(\Theta|D)\right) = \ln\left(p(\Theta)\right) - n\ln\left(\sigma_\epsilon\sqrt{2\pi}\right) - \frac{1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2 \\ - \ln\left(\sum_i p(\Theta_i)p(D|\Theta_i)\right)\end{aligned} \tag{4.16}$$

The last term of the right hand side of Equation 4.16 ultimately ends up being a constant number subtracted from each individual $\ln(p(\Theta_i)p(D|\Theta_i))$ value. Since the value of this constant term does

not impact the shape or relative information of the posterior it could be though of as an arbitrary constant C.

$$\ln\left(p(\Theta|D)\right) = \ln\left(p(\Theta)\right) - n\ln\left(\sigma_\epsilon\sqrt{2\pi}\right) - \frac{1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2 + C \tag{4.17}$$

The constant term can be moved to the left hand side of the equation, producing Equation 4.18.

$$\ln\left(p(\Theta|D)\right) - C = \ln\left(p(\Theta)\right) - n\ln\left(\sigma_\epsilon\sqrt{2\pi}\right) - \frac{1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2 \tag{4.18}$$

Since the objective is to avoid numerical underflow or overflow, prescribing

$$C = \max\left(\ln\left(p(\Theta)\right) - n\ln\left(\sigma_\epsilon\sqrt{2\pi}\right) - \frac{1}{2\sigma_\epsilon^2}\sum_{i=1}^{n}\left(D_i - M_i\right)^2\right) \tag{4.19}$$

shifts all points so that the maximum is 0. A maximum value of 0 in the ln space ensures that all values will be mapped to the interval $[0, 1]$ when taking the exponential. After taking exponentials, the values can be scaled by a constant so that probabilities sum to unity.

4.3.3 Evaluation of Identification Methods

It was desired to evaluate each method for the parameter identification needs of this research. The Bayesian and least squares parameter estimation methods were tested on the five parameter RC network shown in Figure 4.4. In the five parameter RC network heat transfer and storage characteristics of opaque building envelope materials are represented by R_1, R_2, and C. These elements link the ambient temperature node (T_a) to a pseudo interior surface temperature node (T_s), accounting for potential heat storage of the mass materials. Glazing heat transfer is represented by a single resistance (R_w) connecting the ambient temperature node to the surface temperature node, as thermal storage of glazing is typically neglected. R_3 represents a lumped convection/radiation coefficient between the surface temperature node and zone air temperature node (T_z). The convective portion of internal gains (lighting, occupants, and equipment) are applied as a direct heat source to the zone temperature node, shown as $\dot{Q}_{gc}$, and the radiant fraction along with glazing transmitted solar gains ($\dot{Q}_{g,r+sol,w}$) are applied to the surface node.

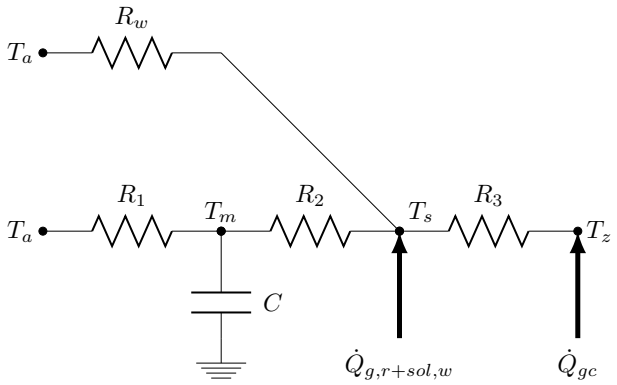

Figure 4.4: Five parameter *RC* network used for evaluation of parameter estimation methods.

Surrogate data from a five zone EnergyPlus retail building was generated for use in the model identification. For the least squares method, 500 direct search points were used within the bounds listed in Table 4.2, and the two-stage least squares parameter estimation was repeated 2500 times beginning each iteration with a new set of randomly generated direct search points. For the Bayesian approach, uniform priors were placed on each resistance and capacitance parameter with the same upper and lower bounds as the least squares method, to express the belief that the best parameters lie within a particular range.

The least squares and Bayesian calibration procedures were repeated considering noise of various colors and amplitudes to better evaluate the ability of each method to handle measurement error and data uncertainty that may occur outside of a simulation testbed. Noise was added to a three week surrogate training data set from the detailed building energy model to simulate sensor error and data uncertainty. *RC* parameters were estimated from the noisy data sets, and performance was compared with respect to a noiseless validation case.

For the least squares method, a data length of one day was also considered to test the impact of significantly reducing the available data. Previous work with inverse gray-box RC models has shown good results using two to three weeks of training data for traditional methods [102, 15]; however, satisfactory performance on shorter data sets would be preferable to reduce data collection burdens. For the Bayesian approach, triangular priors were considered to evaluate the impact of

informative priors. In reality, expert knowledge can often provide more than simple upper and lower bounds. Since Bayesian methods provide a direct means of incorporating such information it was desired to consider this feature in the analysis. White and pink noise were generated using available statistical packages [71, 84, 67], and brown noise for each time step was simulated by adding normal random variations to the noise value of the previous time step as described by Halley and Kunin [39]. Ten noise levels were evaluated that approximately cover the range of 2% to 90% signal error. Additionally, the Bayesian identification was performed for a situation where σ_ϵ was assumed to be known, and a second scenario where σ_ϵ is estimated from the data. Direct comparison between the least squares and Bayesian methods was performed by plotting the 2500 solutions from the least squares method on top of 2-D contour slices of the posterior distribution.

σ_ϵ Known The first estimation task sought to determine the most probable resistance R and capacitance C values assuming a known measurement error standard deviation (σ_ϵ). For a fixed value of σ_ϵ, simulations were performed for 100 000 parameter samples to achieve the posterior distribution. Figure 4.5 shows the posterior slice for parameters R_1 and R_2. Least squares estimates are near probable posterior regions, with a higher R_1 value and lower R_2 value. An inverse relationship between R_1 and R_2 seems reasonable since the parameters together describe envelope heat transfer characteristics. To maintain the same total resistance, an increase in R_1 would require a decrease in R_2.

Figure 4.6 shows that R_2 and R_3 least squares estimates are also associated with probable posterior regions. A potential inverse relationship is also observed for these parameters with the least squares estimating slightly lower envelope resistance (R_2) and higher convection resistance (R_3). This relationship also seems physically plausible since greater heat transfer to the interior surface node could be compensated with reduced heat transfer to the zone air to produce equivalent overall heat transfer.

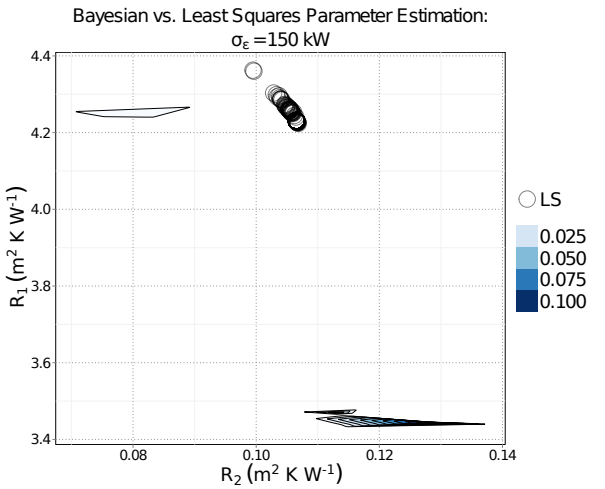

Figure 4.5: Posterior contour for R_1 and R_2 with least squares solutions.

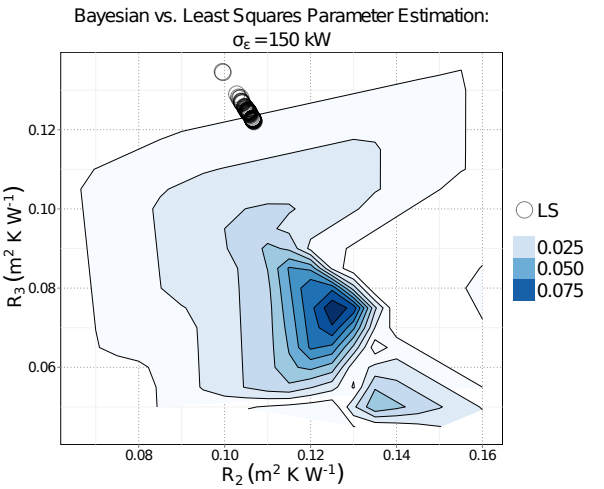

Figure 4.6: Posterior contour for R_2 and R_3 with least squares solutions.

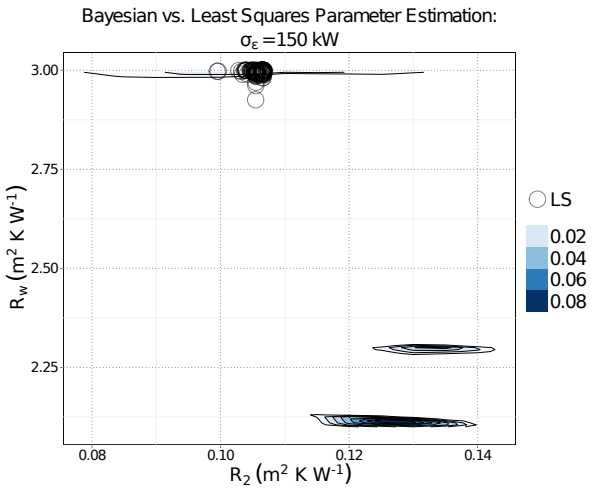

Figure 4.7: Posterior contour for R_2 and R_w with least squares solutions.

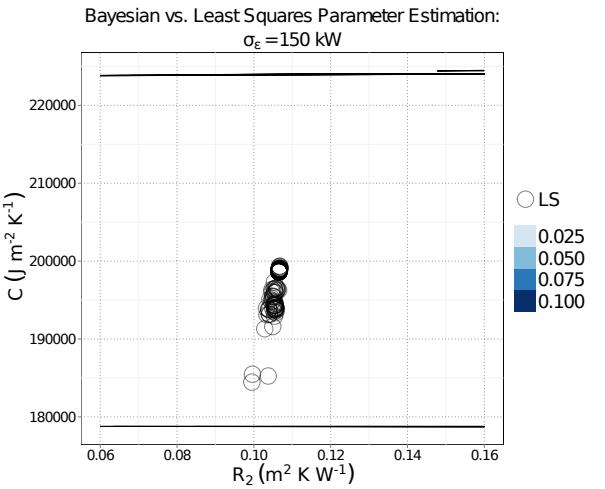

Figure 4.8: Posterior contour for R_2 and C with least squares solutions.

The least squares estimates of R_W were also probable as shown in Figure 4.7. The least squares solutions have higher glazing resistances and slightly lower internal envelope resistance suggesting an inverse relationship between R_w and R_2. Since both parameters are involved in characterizing envelope performance an inverse relationship seems reasonable. However, due to the overlap of probable regions in the R_2 domain, the inverse relationship is not as clearly defined. This may suggest interactions between other parameters as well.

Figure 4.8 highlights the posterior distribution for thermal mass element C, and shows that least squares estimates were spread between two probable regions. Although seemingly different, the range of capacitance between these two regions is approximately equivalent to a concrete wall of thickness ranging from 9.6 cm to 12 cm (3.78 to 4.72 in.), which represent fairly similar constructions. Parameter bounds (priors) were set to encapsulate a large range of construction possibilities, however this range is numerically much larger for the C parameter. A further investigation of the results showed that only 8% of the sampled C values were in the range between two probable regions. The wide parameter bounds (priors) may have prohibited adequate sampling to provide high resolution posterior information in this region. The results could potentially be improved by using the initial estimates to update the parameter bounds and narrow the solution space. Similar results are observed in the R_3C, R_wC, and R_1C posterior slices in Figs. 4.9, 4.10, and 4.11 respectively.

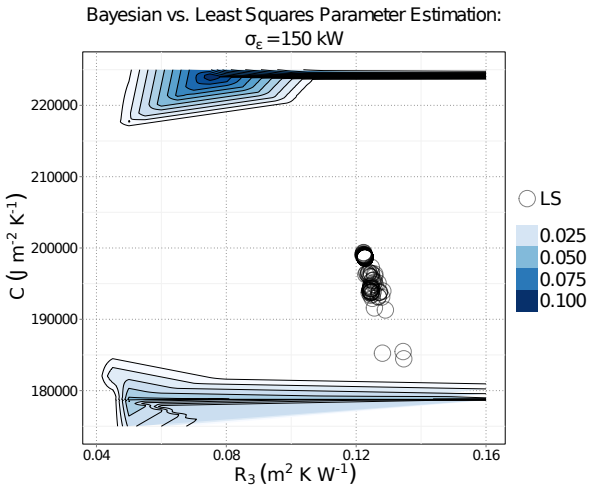

Figure 4.9: Posterior contour for R_3 and C with least squares solutions.

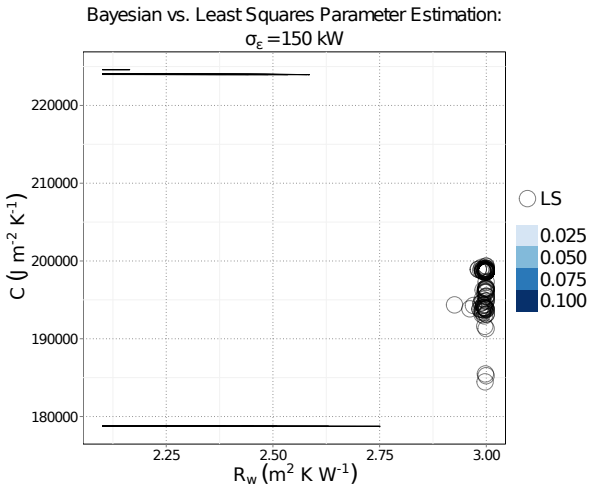

Figure 4.10: Posterior contour for R_w and C with least squares solutions.

Potential direct interactions between the glazing resistance (R_w), external envelope resistance (R_1), and the internal convection resistance (R_3) are shown in Figure 4.12 and Figure 4.13. This relationship may also have a physical explanation since decreased envelop heat transfer could result

in higher internal surface node temperatures. To maintain the same overall heat transfer to the zone node an increase in the convection resistance would be required.

Parameter estimates are summarized in Table 4.3, and the largest discrepancy between the two methods was observed in the external envelope resistance (R_1) and window resistance (R_w) estimates shown visually in Figure 4.14. This difference may suggest that the model is relatively insensitive to parameters R_1 and R_w, and consequently insensitive to ambient temperatures. This interpretation seems plausible since heat transfer in commercial buildings is often dominated by internal gains, and this particular building has notably high lighting loads.

Table 4.3: Retail model parameter estimates.

	NLSQ		Bayes: σ_ϵ assumed		Bayes: σ_ϵ estimated		
	median	95% Conf.	p_{max}	95% Cred.	p_{max}	95% Cred.	Units
R1	4.231	[4.231, 4.231]	3.445	[0.385, 4.876]	4.306	[4.041, 4.585]	m^2 K W^{-1}
R2	0.107	[0.107, 0.107]	0.128	[0.013, 4.409]	0.111	[0.096, 0.120]	m^2 K W^{-1}
R3	0.123	[0.123, 0.123]	0.076	[0.048, 31.98]	0.113	[0.099, 0.142]	m^2 K W^{-1}
Rw	2.999	[2.999, 2.999]	2.114	[0.214, 2.999]	2.078	[1.728, 2.837]	m^2 K W^{-1}
C	198.7	[198.7, 198.7]	224.0	[30.80, 511.6]	200.1	[184.3, 230.3]	kJ m^{-2} K^{-1}
σ_ϵ	-	-	150.0	-	11.00	[10.26, 11.30]	kW

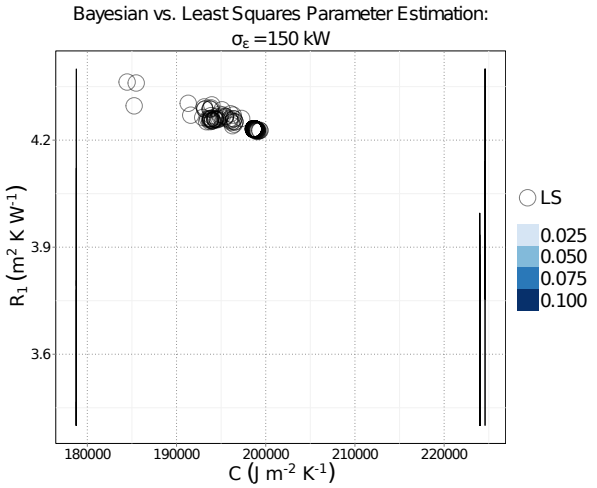

Figure 4.11: Posterior contour for C and R_1 with least squares solutions.

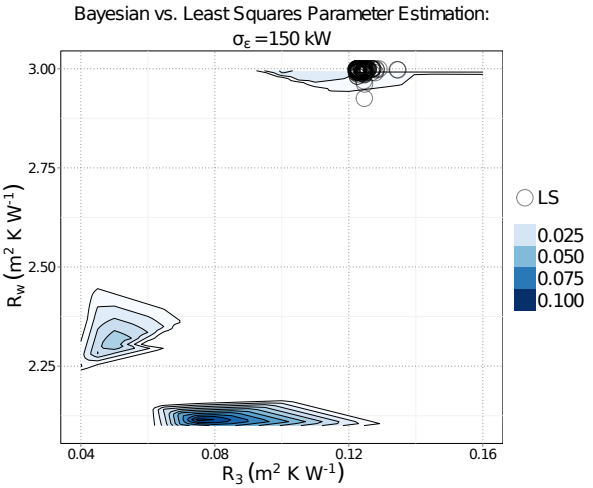

Figure 4.12: Posterior contour for R_3 and R_w with least squares solutions.

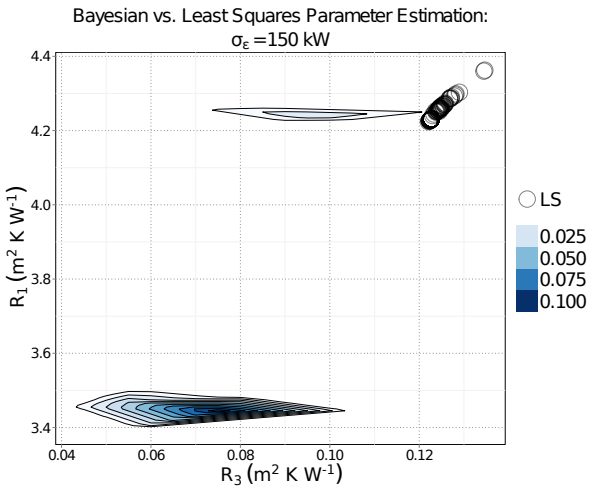

Figure 4.13: Posterior contour for R_3 and R_1 with least squares solutions.

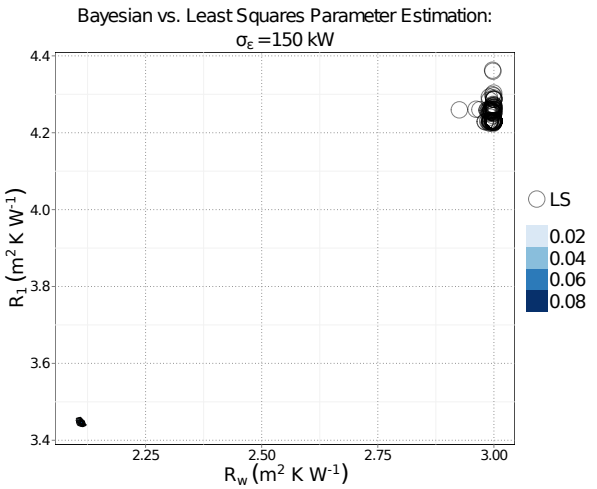

Figure 4.14: Posterior contour for R_w and R_1 with least squares solutions.

Sensitivity to σ_ϵ The maximum posterior probability in the previous results was 0.1. While seemingly low, the magnitude of these probabilities is inherently related to the uncertainty in the problem, namely the measurement error, and the fact a time series of data are considered for the parameter estimation task. Looking back to Equation 4.17 shows the posterior probability is a function of the number of time series data points n and the measurement error σ_ϵ, along with posterior probabilities and the sum of squared errors. If the measurement error is assumed too large all parameter combinations are similarly likely; too small and no parameters are likely, resulting in a meaningless posterior distribution.

For the results above, σ_ϵ was initially assumed to be 150 kW (about 50% peak heating load) in order to account for potentially large deviations that could result from model simplification and abstraction. An error standard deviation of 150 kW was found to clearly distinguish the best parameters while still providing useful information about other parameter combinations and potential interactions.

The posterior was re-computed for various values of σ_ϵ to explore how the posterior changes with measurement error. Figure 4.15 shows how the maximum posterior probability changes with

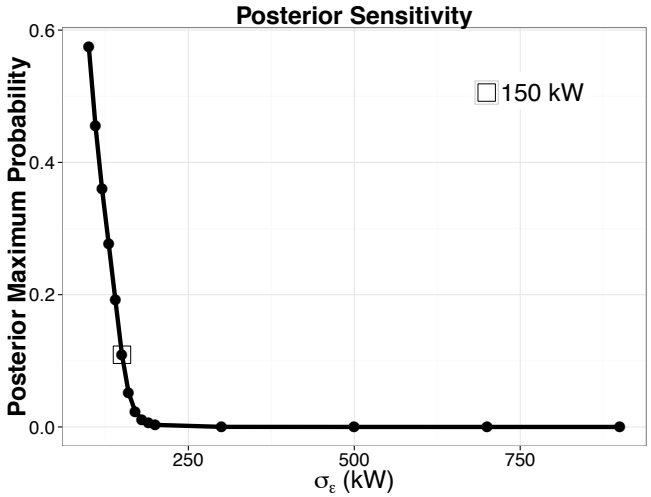

Figure 4.15: Posterior maximum for various values of σ_ϵ.

σ_ϵ. As σ_ϵ is increased from 150 kW the posterior maximum decreases since greater error allows more parameter sets to have non-zero probability. Similarly, as σ_ϵ decreases from 150 kW, the posterior maximum increases as probabilities are concentrated among fewer parameter combinations. These results also suggest that the best estimate of σ_ϵ may still be lower than 100 kW (about 33% max heating load) since the point where no parameter combinations are likely has not yet been determined.

σ_ϵ Estimated Since the value of σ_ϵ is typically not known, a natural extension is to estimate the measurement error standard deviation along with model parameters. The Bayesian inference was repeated with σ_ϵ treated as a free (sixth) parameter to estimate the most likely value. A uniform prior from 0.1 kW to 250 kW (about 0.03% – 80% max heating load) was place on σ_ϵ, which was sampled along with R and C model parameters. The number of samples was also increased to 500 000 to account for the increase in dimensionality. The posterior maximum for all parameters occurred with a σ_ϵ value of 11 kW (about 4% max heating load), and as Figure 4.15 alluded, was much lower than the initially assumed 150 kW.

In general, the most probable parameter set from the six parameter estimation were in agreement with the least squares solutions. Figures 4.16 – 4.25 show the updated posterior slices

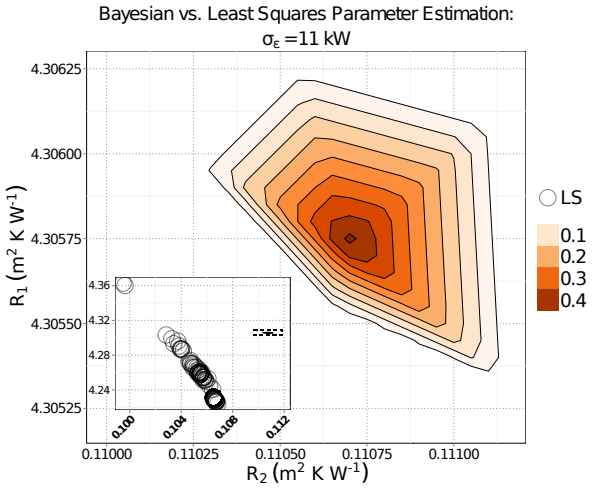

Figure 4.16: Posterior contour for R_2 and R_1 with least squares solutions.

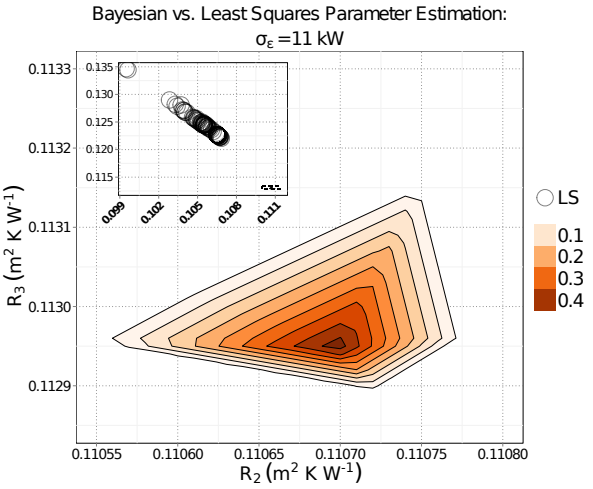

Figure 4.17: Posterior contour for R_2 and R_3 with least squares solutions.

when σ_ϵ is included in the estimation.

Figure 4.16 shows that several least squares runs were in agreement with the Bayesian inference for R_1, although the median value is slightly lower. R_2 values were in close agreement as well, although slightly higher for the Bayesian results. Since R_1 and R_2 represent the envelope material resistances, choosing higher values for both parameters is not directly explainable from Figure 4.16. However, looking ahead, the Bayesian inference chose a lower glazing resistance, which may result in similar overall envelope performance.

Figure 4.17 again shows slightly higher values for the Bayesian estimate of R_2 and lower values for the internal convection resistance R_3. This inverse relationship between envelope resistance and internal convection resistance may result in similar heat transfer to the zone air node.

Figure 4.18 directly highlights the inverse relationship previously discussed between R_w and R_2 with the Bayesian inference choosing reduced glazing resistance in favor of increased external wall resistance.

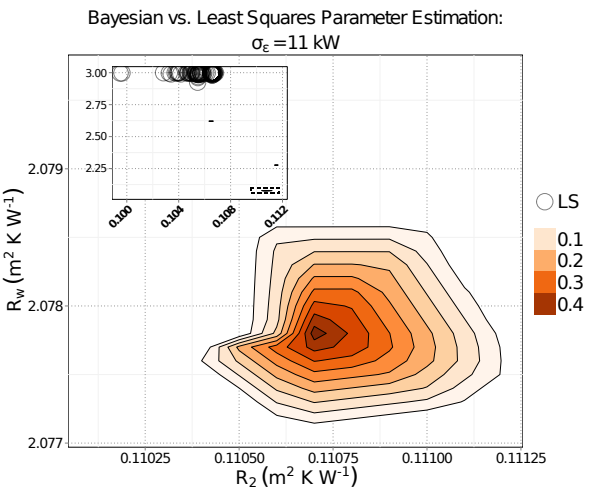

Figure 4.18: Posterior contour for R_2 and R_w with least squares solutions.

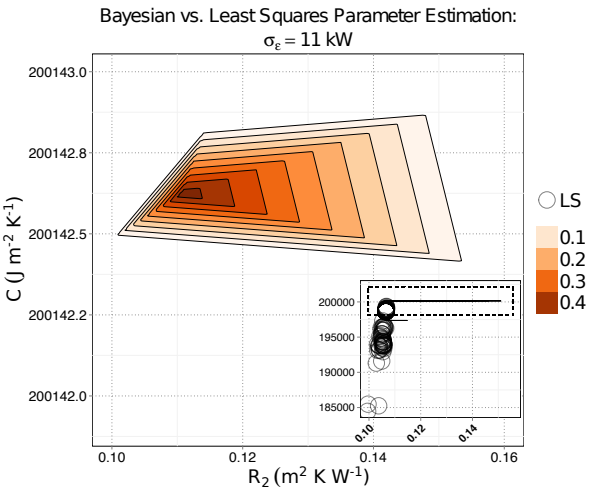

Figure 4.19: Posterior contour for R_2 and C with least squares solutions.

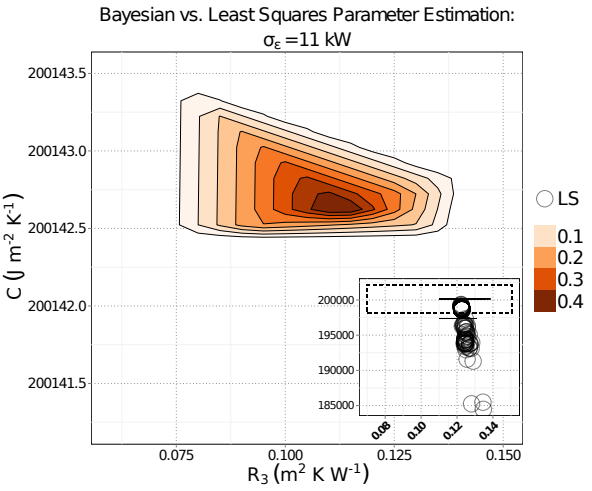

Figure 4.20: Posterior contour for R_3 and C with least squares solutions.

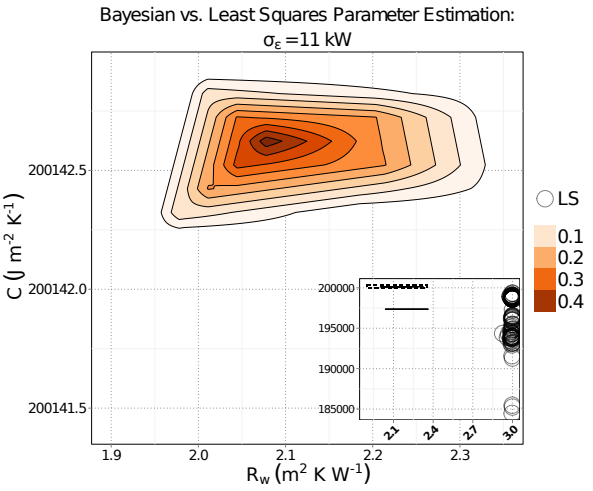

Figure 4.21: Posterior contour for R_w and C with least squares solutions.

Figures 4.19 – 4.22 show the close agreement between thermal capacitance values (C). Although no least squares points appear directly on the posterior, the difference between chosen C values is less than 1%. The remaining parameter combinations are plotted in Figures 4.23 – 4.25, and parameter values are summarized in Table 4.3.

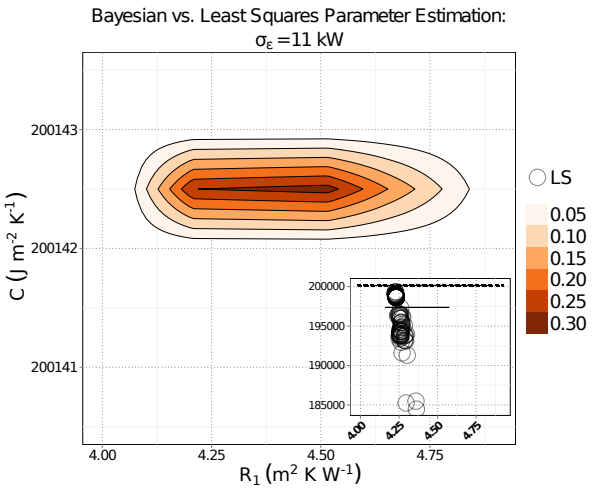

Figure 4.22: Posterior contour for C and R_1 with least squares solutions.

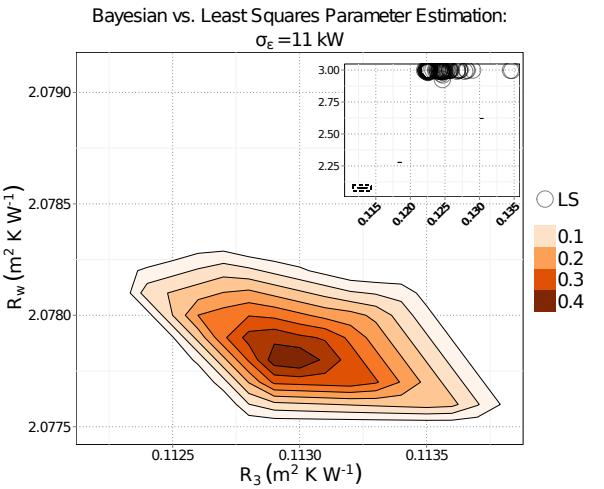

Figure 4.23: Posterior contour for R_3 and R_w with least squares solutions.

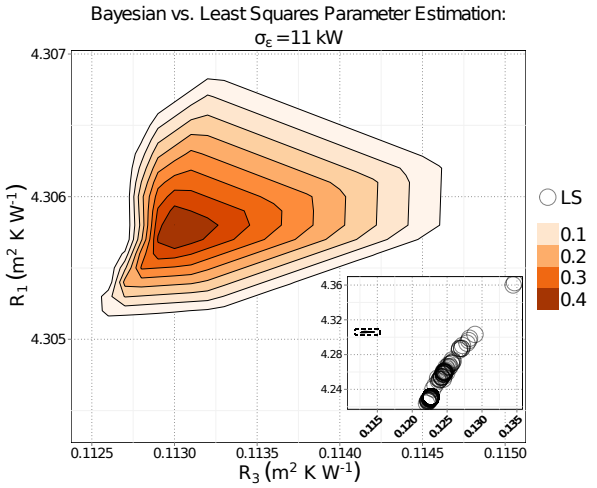

Figure 4.24: Posterior contour for R_3 and R_1 with least squares solutions.

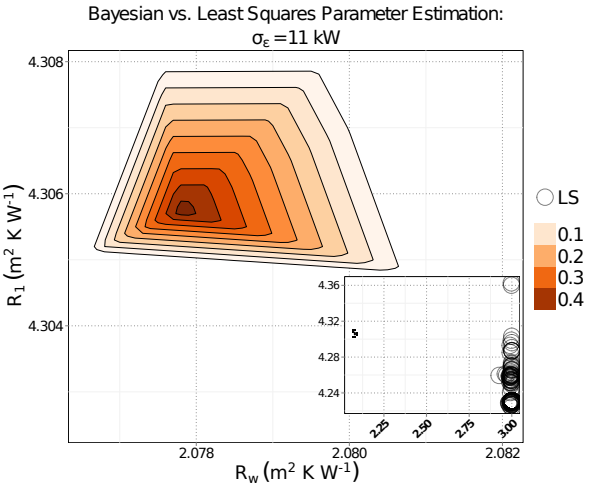

Figure 4.25: Posterior contour for R_w and R_1 with least squares solutions.

Parameter Performance Parameter estimates were simulated for a validation week to compare predicted performance. All three parameter sets were able to adequately forecast the heating loads and zone temperatures as shown in Figure 4.26 and error metrics in Table 4.4, where $RMSE$ refers to root mean square error and MBE to mean bias error. However, it should be

noted that from the Bayesian perspective distance metrics are meaningless and every inference is phrased in terms of probability.

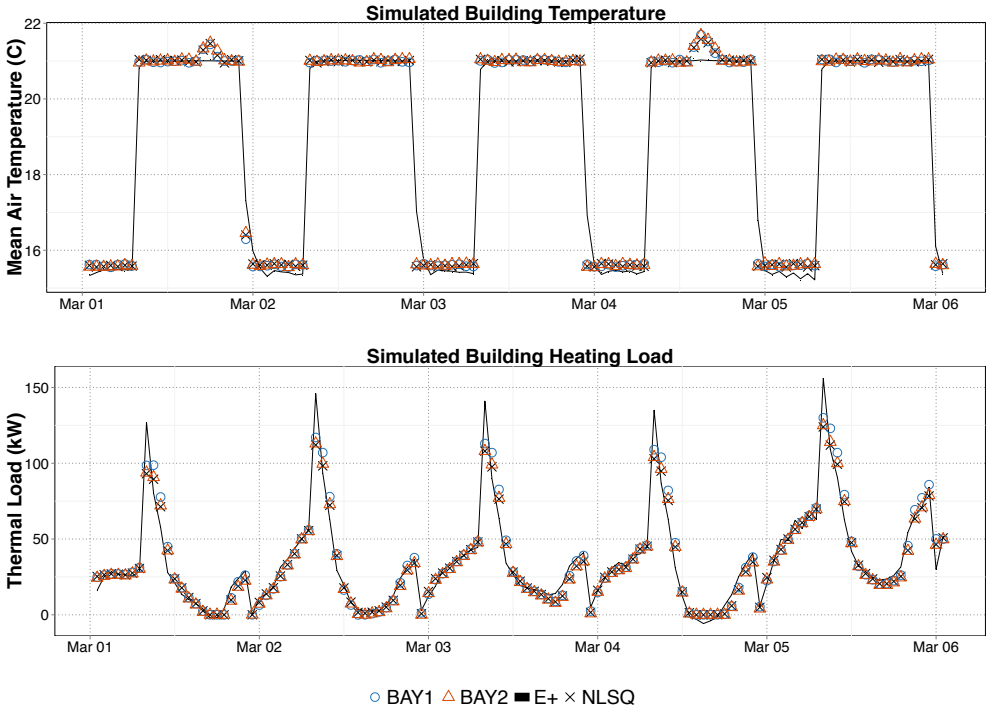

Figure 4.26: Performance comparison of Bayesian and least squares estimates.

Table 4.4: Model performance summary metrics.

		RMSE	MBE	Cum. % Err.
Load	NLSQ	10 250 W	347 W	0.869%
	Bayesian (σ_ϵ assumed)	10 692 W	1802 W	4.491%
	Bayesian (σ_ϵ estimated)	10 244 W	-24 W	-0.059%
Temp.	NLSQ	0.2974 K	0.0360 K	0.191%
	Bayesian (σ_ϵ assumed)	0.2999 K	0.0367 K	0.195%
	Bayesian (σ_ϵ estimated)	0.2991 K	0.0386 K	0.205%

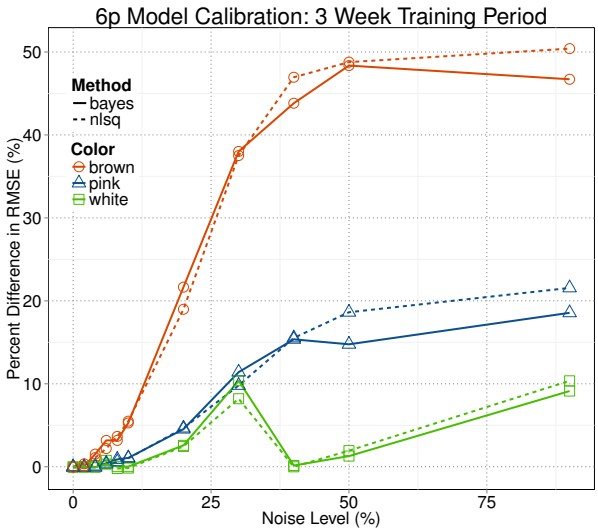

Figure 4.27: RMSE percent difference with three week training period and uniform priors.

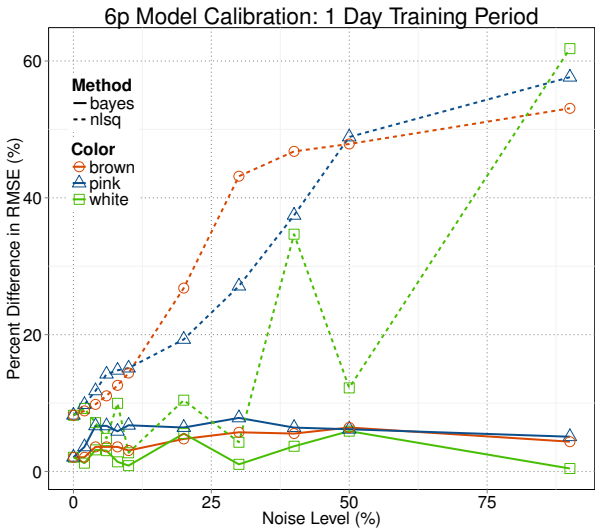

Figure 4.28: RMSE percent difference with 24 hour training period and triangular priors.

Sensitivity to Noise According to the experiment plan, the estimation was repeated after adding noise to the three week training period to explore the sensitivity of each method to various levels of noise and noise color. Colored noise is a proxy for correlated and uncertain phenomena, and it is important to ensure the thermal model calibration is robust in the face of real measured data, especially where load is not known but estimated.

The Bayesian estimation was performed including estimation of σ_ϵ with uniform priors on all parameters. Performance was evaluated by computing $RMSE$ on a noiseless validation period, and the results are compared by normalizing with respect to the least squares case trained on three weeks of noiseless data (i.e. RMSE reported in Table 4.4). Figure 4.27 summarizes the results and shows that both methods perform similarly for all noise levels and colors tested (i.e. marker shapes group together). It should be noted that, for a given data set, similar performance is expected when using uniform priors since all parameters are equally likely between bounds and the likelihood is based on the sum of squared deviations (i.e. finding maximum posterior probability parameters is equivalent to finding least squares parameters in this scenario).

To consider the impact of nonuniform priors, the noise analysis was repeated using triangular prior distributions on the parameters. A truncated data set of 24 hours was used to evaluate the change in least squares performance due to data availability as well. It is noted that least squares results in Figure 4.27 and Figure 4.28 can be compared directly since they vary only by data length; however, Bayesian results can not directly be compared since they differ by both data length and prior information. The results have been normalized with respect to the least squares reference case trained on three weeks of noiseless data.

Figure 4.28 shows that the methods perform similarly for low noise levels and that the Bayesian approach outperforms for higher noise levels (i.e. line types group together). The improved performance of Bayesian methods at higher noise levels suggests that the use of informative priors can buffer potentially misleading information present in noisy data sets.

Comparing least squares results between Figure 4.27 and Figure 4.28 shows that performance is lost when less data is available, however for most scenarios only 5 – 10% increase in RMSE was observed. It is also interesting to highlight that the Bayesian calibration acknowledges the increasing levels of data noise through higher estimates of σ_ϵ as shown in Figure 4.29.

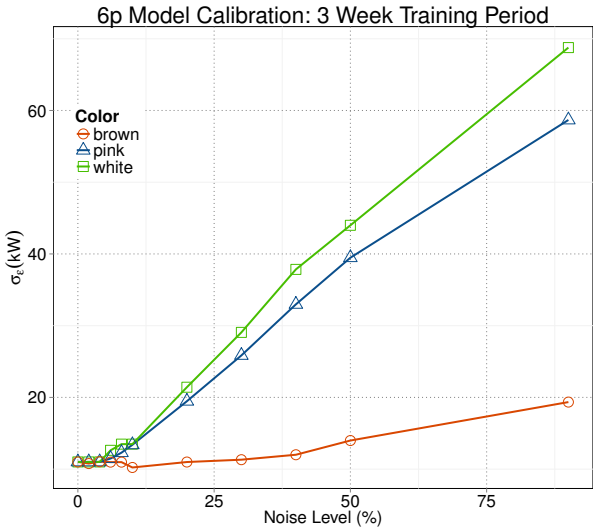

Figure 4.29: Measurement error standard deviation as a function of noise level.

Summary of Parameter Estimation Evaluation Least squares and Bayesian parameter estimation methods performed similarly, regardless of noise, when using uniform priors. The Bayesian methods did show the potential to outperform on noisy (i.e. $> 25\%$ noise level) data sets when utilizing informative priors. Bayesian estimation was also able to provide further insight into potential parameter interactions and trade-offs, as well as parameter uncertainty. However, the additional information does not come without added computational cost. The Bayesian methods performed in this work required approximately 100 times more CPU time than the traditional parameter estimation. The initial objectives of this research sought to evaluate the opportunity for synergistic effect among building portfolios in a deterministic testbed; and, therefore, do not truly benefit from the probabilistic Bayesian perspective. In an online operational setting, the Bayesian methods may be preferred if uncertainty quantification is desired or high levels of sensor noise are expected. Furthermore Bayesian methods may also show additional benefit on more complex, higher-dimensional models with multiple sources of noise and uncertainty.

4.4 Model Complexity

When utilizing the inverse gray-box modeling approach described in this work, questions naturally arise around what RC network structure is most appropriate for the modeling task. Selecting a very complex model structure results in difficult parameter estimation, while too simple of a model may not appropriately capture the desired dynamics. In this research the reduced-order modeling environment has been developed to allow for model structure flexibility so this question may be investigated. As previously mentioned, various RC network forms have been considered in this work ranging from five to 21 parameters. Since the 21 parameter model has been previously introduced, discussion will begin with the 18 parameter model.

The 18 parameter model, shown in Figure 4.30, can be considered a subset of the 21 parameter model with the internal surface heat transfer elements simplified to 1R1C. This reduced the parameter estimation procedure by three parameters while keeping the majority of the structure of the 21 parameter model. The 13 parameter model, shown in Figure 4.31, is a subset of the

21 parameter model as well with a simplified internal surface node and no ground heat transfer. The initial concept for this model is that for small footprint high-rise buildings the ground heat transfer may not be a significant contributor to the overall thermal load. The eleven parameter model (Figure 4.32) contains the simplified internal surface network of the 18 parameter model, as well as a simplified ground heat transfer network and lumped ceiling and exterior wall networks. The eight parameter model, shown in Figure 4.33 further simplifies the eleven parameter model by neglecting ground heat transfer. This model contains a 3R2C network for exterior surfaces, a glazing resistance, and a simplified internal surface/mass network. The five parameter model, shown in Figure 4.34 is based on a structure independent of the original 21 parameter model. The five parameter structure was adapted from the thermal RC network used in the ISO 13790 "Simple Hourly Method" load calculations discussed in Section 4.1, and was previously described in further detail in Section 4.3.3.

As previously mentioned it is of interest to evaluate the performance of various forms of thermal RC networks. An initial hypothesis was that more complex buildings may benefit from more complex thermal RC networks. To investigate this theory, each of the six models (5p, 8p, 11p, 13p, 18p, 21p) were trained for each of the building applications (stand-alone retail, medium office, and large office). The retail model was trained using surrogate data from a five zone DOE Stand-alone Retail Reference EnergyPlus model, the medium office was trained from surrogate data from a six-story, 15 zone EnergyPlus office building model, and the large office model was trained from a 30-story, 15 zone EnergyPlus office building model. Table 4.5 summarizes the model performance in terms of RMSE with respect to a validation data set, as well as in terms of an objective generalized cross-validation score (GCV). GCV is defined in Equation 4.20 and essentially weights the mean-squared error based on model complexity [11].

$$GCV = \frac{\sum_{i=1}^{N} (\dot{Q}_{rom,i} - \dot{Q}_{ep,i})^2}{N\left(1 - \frac{p}{N}\right)^2} \tag{4.20}$$

In Equation 4.20, N represents the total number of data points, $\dot{Q}_{rom,i}$ is the model predicted zone sensible load, $\dot{Q}_{ep,i}$ is the surrogate zone load, and p is the number of parameters in the model.

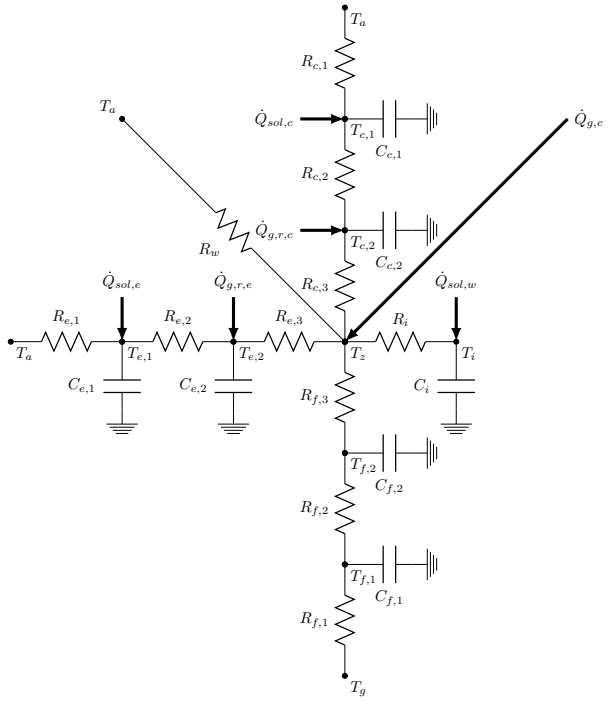

Figure 4.30: 18 parameter thermal RC network.

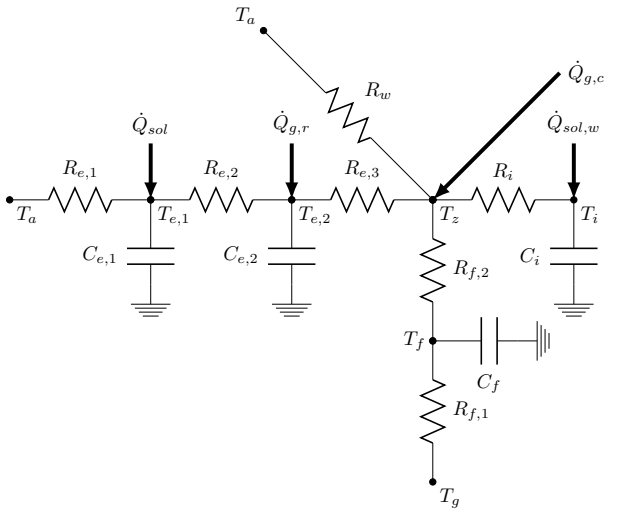

Figure 4.32: Eleven parameter thermal RC network.

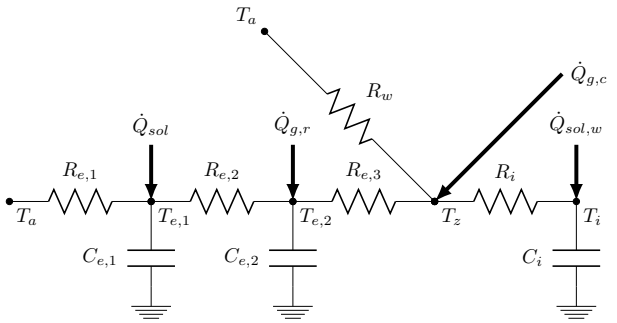

Figure 4.33: Eight parameter thermal RC network.

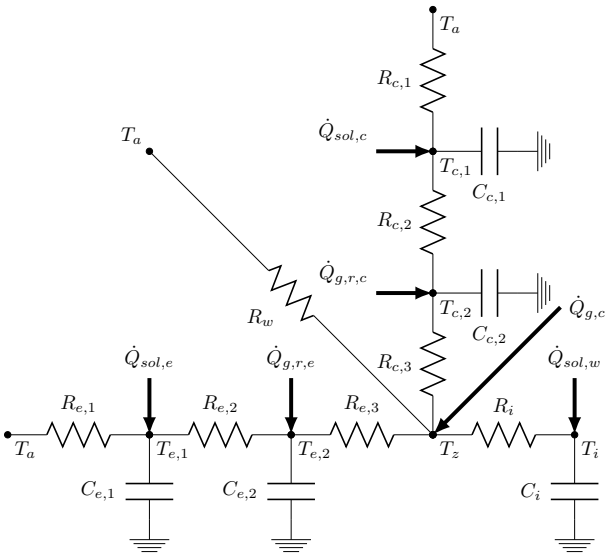

Figure 4.31: 13 parameter thermal RC network.

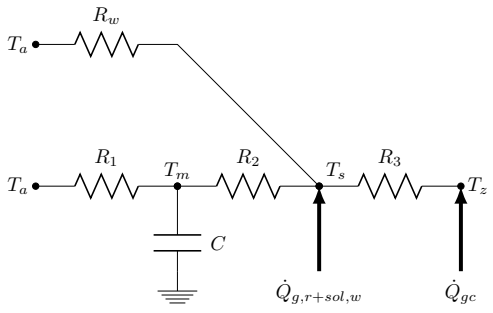

Figure 4.34: Five parameter RC network.

Table 4.5: Model complexity results.

				Retail		Medium Office		Large Office	
Model	p	N	k	RMSE	GCV	RMSE	GCV	RMSE	GCV
5p	7	504	128	6411	42.3×10^6	55 520	3.2×10^9	635 398	415.2×10^9
8p	10	504	1024	5338	29.7×10^6	48 967	2.5×10^9	278 207	80.6×10^9
11p	13	504	8192	3087	10.0×10^6	30 492	1.0×10^9	252 262	67.1×10^9
13p	15	504	32768	5234	29.1×10^6	51 445	2.8×10^9	269 681	77.3×10^9
18p	20	504	1048576	3076	10.3×10^6	31 885	1.1×10^9	261 980	74.4×10^9
21p	23	504	8388608	3192*	11.2×10^6	26 967	0.8×10^9	162 982	29.2×10^9

*Slightly suboptimal. Should have at least reach 3076 as the 18p retail model.

The number of parameters p is equal to the number of RC parameters plus two, to account for the internal gain and zone capacitance multipliers. The model complexity analysis was performed using the least squares parameter identification approach with the RMSE objective function previously defined by Equation 4.5.

Table 4.5 shows that for the large office building the 21 parameter model resulted in the lowest RMSE and lowest GCV. For the medium office, the 21 parameter model also achieved the lowest RMSE and GCV, although the eleven parameter performed similarly as well. For the retail building the 18 parameter model produced the lowest RMSE, however the 21 parameter results are in adequate since they should have at least achieved the same score as the lower order model. However the 11p model produced the lowest GCV, and suggests that the additional improvement made by the 18p model was not worth the additional complexity. The RMSE values for the retail building are all relatively low, also suggesting that all model forms performed well. In general, the five parameter model resulted in the highest RMSE and GCV values for all building types.

An interesting facet is observed when comparing models in the middle of the complexity range. RMSE typically increases as the model is simplified from 21 to thirteen parameters. The eleven parameter model often performs similar to or better than the 18 parameter model, and the RMSE once again increases as the parameters are reduced towards the five parameter model. A further investigation of this observation revealed that model simplification from 21 to five parameters was

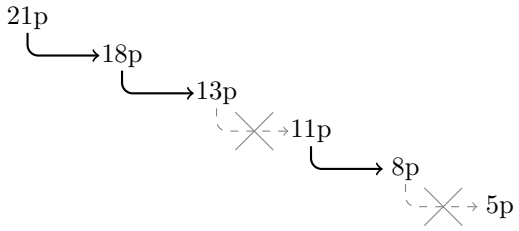

Figure 4.35: RC model simplification hierarchy.

not continuous. Figure 4.35 highlights that the 13 parameter model is unable to simplify to the eleven parameter model, and the eight parameter model cannot simplify to the five parameter model for the constrained optimizations performed in this work.

This discontinuous simplification can be described by comparing the eleven and 13 parameter models in Figures 4.32 and 4.31, respectively. The differences in model structure lie in the eleven parameter model having a 2R1C network connected to a ground temperature node, while the thirteen parameter instead has separate 3R2C network describing the roof heat transfer. The ground network in the 11p model could be reduced to match the 13p model by choosing high resistances and insulating the effects of the ground temperature forcing function. However, the 3R2C roof network cannot be simplified in the 13p model by similar logic. Simply choosing a high resistance for $R_{c,3}$ would result in a significant reduction in the internal gains since the effects of the $\dot{Q}_{g,r,c}$ forcing function would be eliminated. In an unconstrained optimization the 13p model should, technically, be able to collapse the roof and envelope networks to reach the 11p envelope structure, since they are parallel branches. However, the parameter constraints used in these identification tasks limited the opportunities for simplification to a very small parameter set. These parameters did not lead to satisfactory model results and were not chosen by the optimizer for any of the buildings. The eight parameter model cannot simplify to the five parameter model due to significant differences in structure and location of applied heat sources.

4.5 HVAC System Modeling

The following sections provide an overview of the systems modeled for each building application. Detailed system and component validation results are provided in Appendix B. In general HVAC component models were modeled after similar methods presented in the EnergyPlus Engineering Reference manual as well as the ASHRAE HVAC Toolkit 2 [13]. Component models were programmed such that a full air loop can be simulated. Full air loop simulation allows system air states to be included in a simple zone moisture balance so hourly zone humidity estimates can be provided.

4.5.1 Packaged DX RTU

For the stand alone retail building a typical constant volume packaged roof-top unit was modeled. Figure 4.36 provides an overview of the system configuration. The RTU model features a temperature or enthalpy based outdoor air economizer, constant volume fan, single speed DX cooling coil, and gas heating coil.

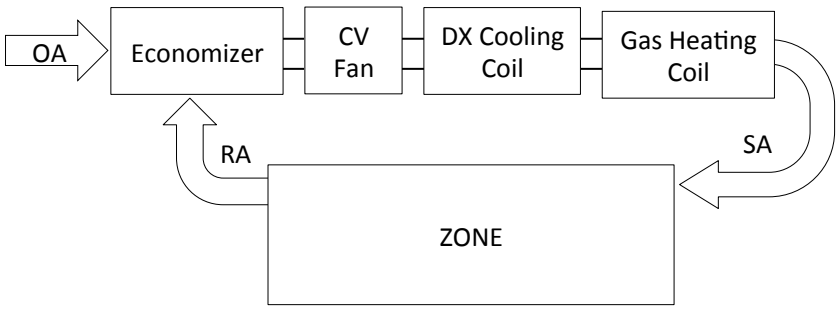

Figure 4.36: Retail building packaged DX roof top unit model.

4.5.2 Packaged DX VAV

For the medium office building a packaged VAV system was modeled. The system incorporates a dedicated outdoor air system (DOAS) with supply and exhaust fans, gas heating coil, DX cooling coil, and variable speed fan. The original EnergyPlus model contains a hot water heating coil with

hot water reheat terminal units in the perimeter. This system approximates the AHU heating coil with a gas coil and terminal reheat as a unit heater. It is apparent in the validation results that making the single zone assumption results in the inability to accurately predict winter reheat loads. However, since this work focused on cooling season performance these simplifications were adequate.

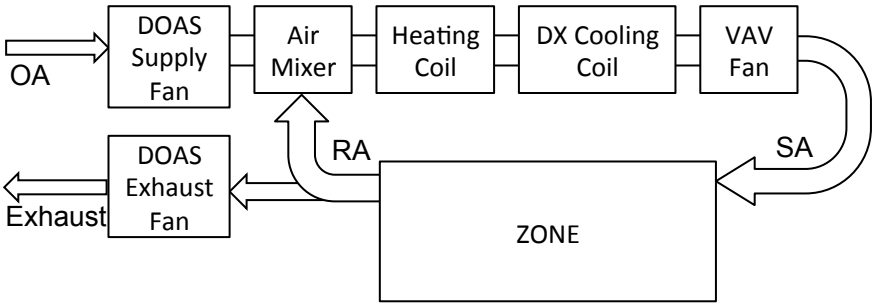

Figure 4.37: Medium office packaged DX VAV system model.

4.5.3 Chilled Water VAV

Figure 4.38 shows the chilled water VAV system modeled for the large office building. The air handler model incorporates an enthalpy based outdoor air economizer, VAV supply and return fans, and chilled water cooling coil. The cooling plant consists of two electric chillers, two cooling towers, headered condenser water pumps, and headered chilled water pumps.

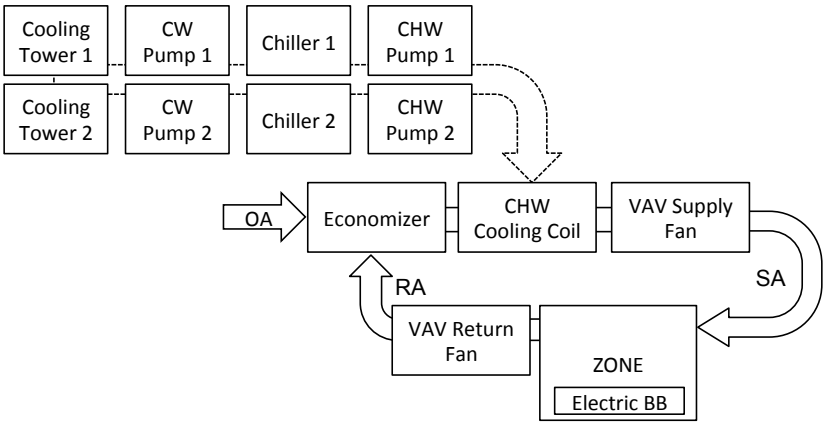

Figure 4.38: Large office chilled water VAV system model.

4.6 Whole Building Model Validation

4.6.1 Retail Building

The following results provide a comparison of the overall (i.e. single-zone inverse gray-box building model and MATLAB HVAC models) reduced-order retail model performance compared to its five zone EnergyPlus counterpart. To provide better insight into the model performance under different conditions the model was simulated using typical NSU (Night Set-Up) operation during a mild week and a pre-cooling heuristic for a hot week. (It is noted that these are validation time periods—that is, neither were included in model training.) Zone sensible load, temperature, and HVAC electric consumption are in good agreement for both NSU and precooling scenarios in Figures 4.39 – 4.44.

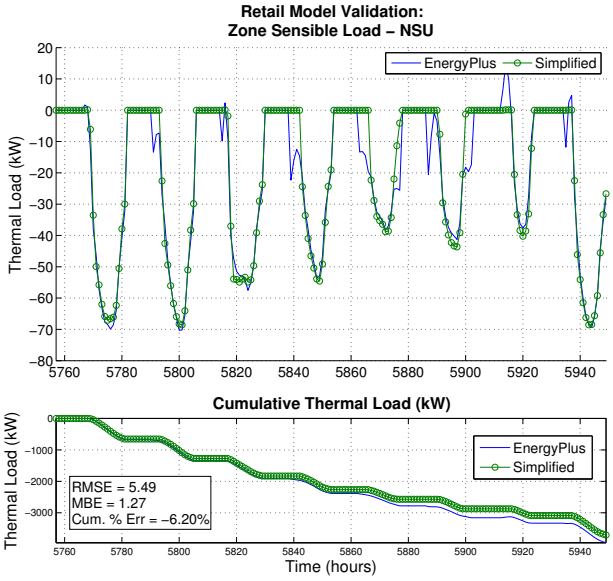

Figure 4.39: Retail sensible zone load comparison for NSU scenario.

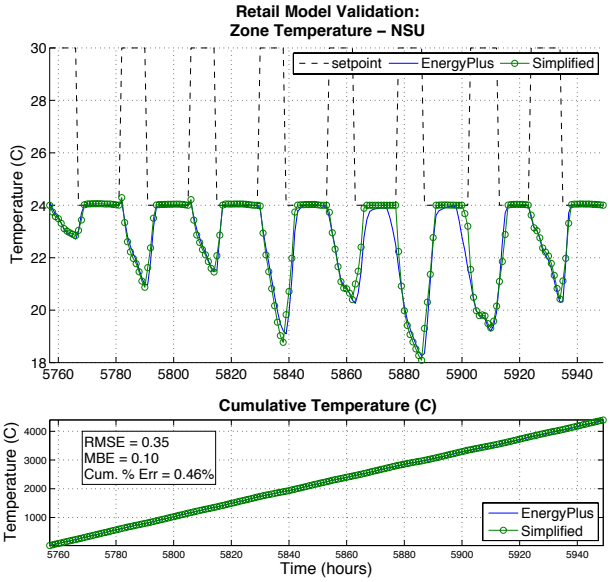

Figure 4.40: Retail zone mean air temperature comparison for NSU scenario.

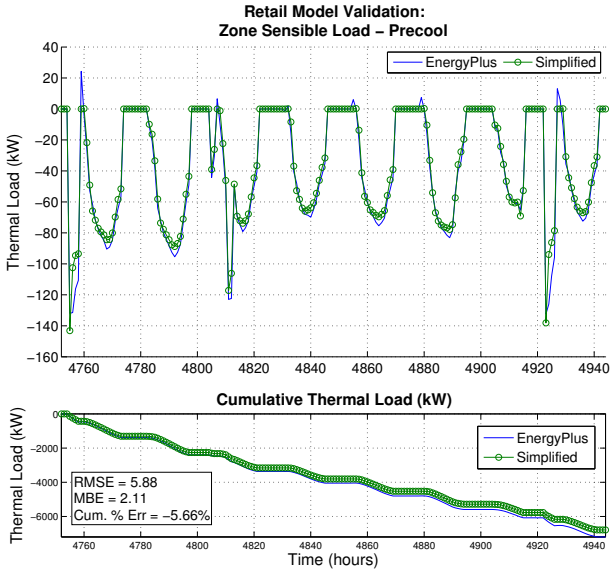

Figure 4.41: Retail sensible zone load comparison for precooling scenario.

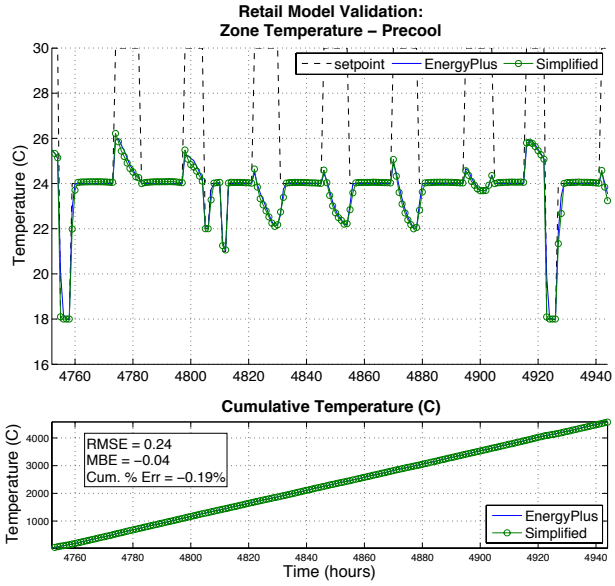

Figure 4.42: Retail zone mean air temperature comparison for precooling scenario.

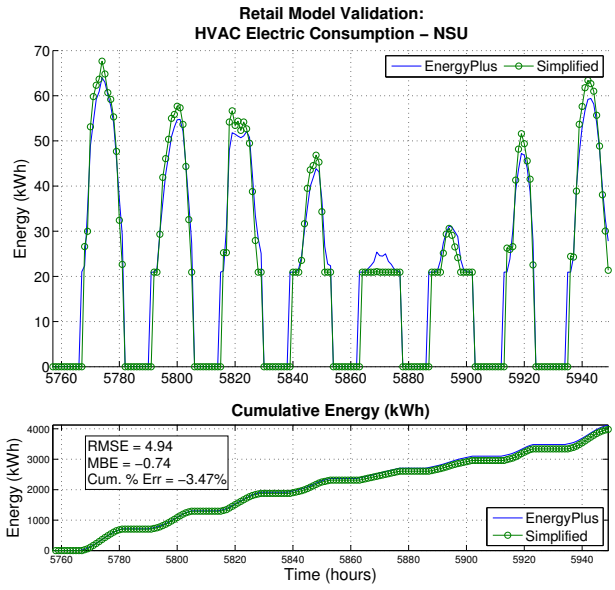

Figure 4.43: Retail HVAC electric consumption comparison for NSU scenario.

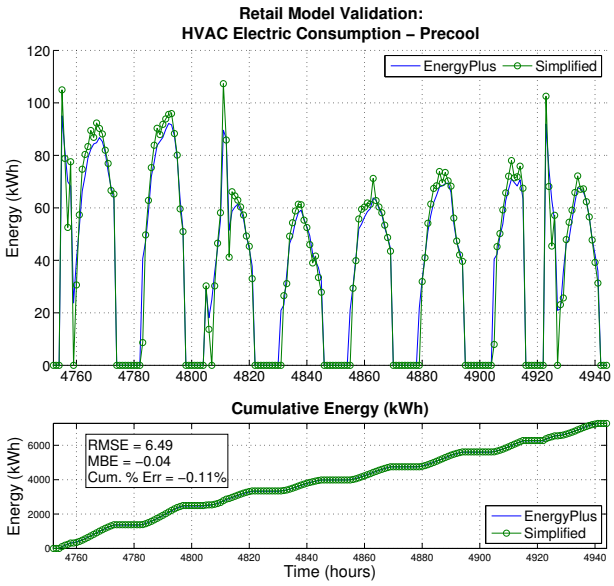

Figure 4.44: Retail HVAC electric consumption comparison for precooling scenario.

4.6.2 Medium Office Model Comparison

A similar comparison was made for the medium office model. Good results were also observed for both NSU and precooling scenarios in Figures 4.45 – 4.50. An important model characteristic is observed during the mild week temperatures in Figure 4.46. Temperatures are below the cooling setpoint for both EnergyPlus and reduced-order models due to minimum VAV air flow limitations. Since no SAT reset is in effect, the minimum airflow rates are slightly overcooling the zone. Reduced-order models are able to capture this since the load calculations are updated based on available HVAC capacity. System capacity limitations can also be observed in Figure 4.48 during deep precooling near hour ending 5260. Data markers are placed at hourly intervals and it can be seen that it takes approximately four hours to precool to 20.5 °C.

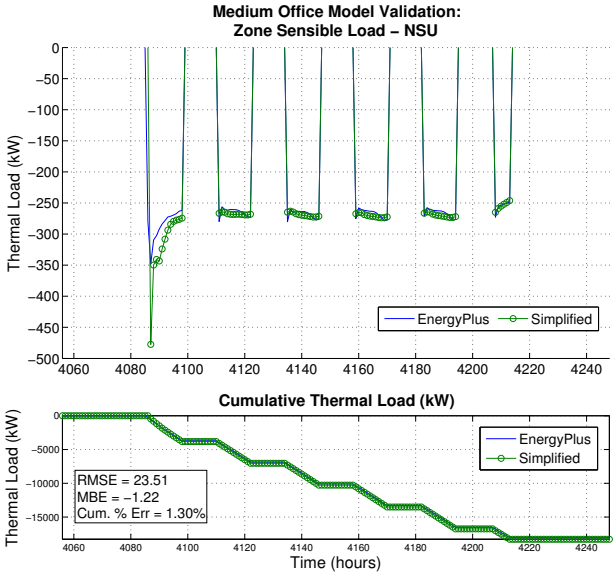

Figure 4.45: Medium office sensible zone load for NSU scenario.

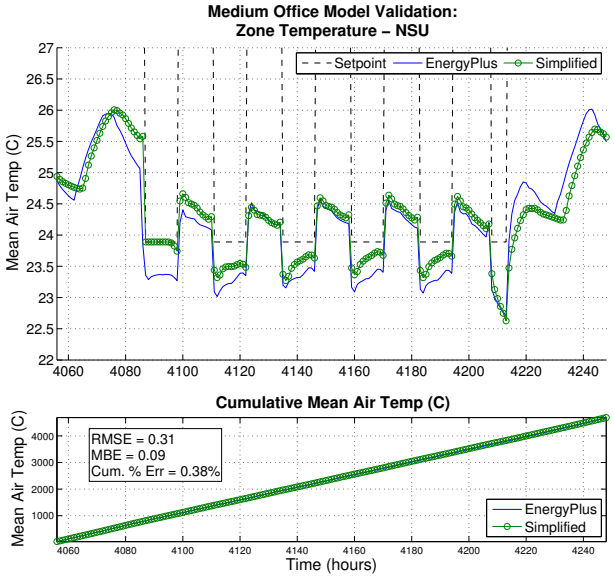

Figure 4.46: Medium office zone mean air temperature for NSU scenario.

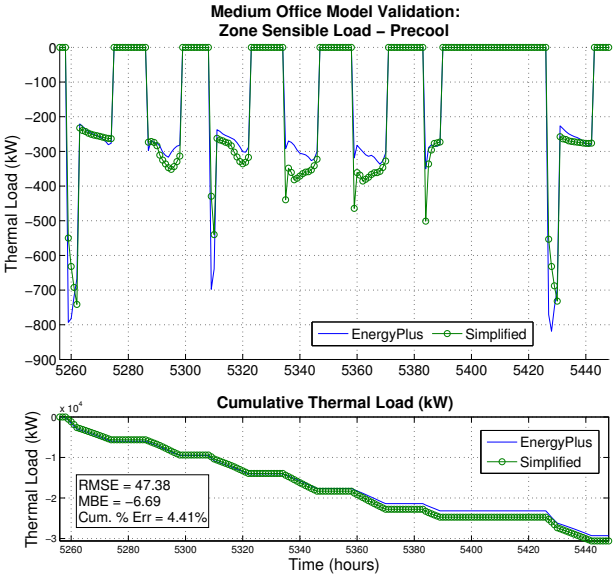

Figure 4.47: Medium office sensible zone load for precooling scenario.

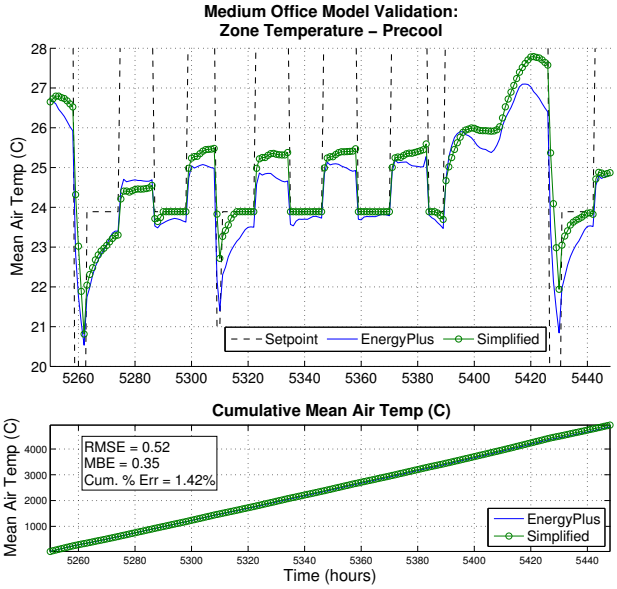

Figure 4.48: Medium office zone mean air temperature for precooling scenario.

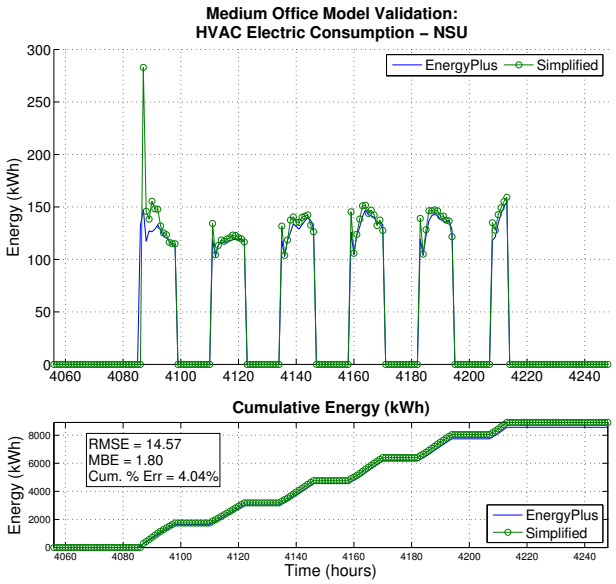

Figure 4.49: Medium office HVAC electric consumption for NSU scenario.

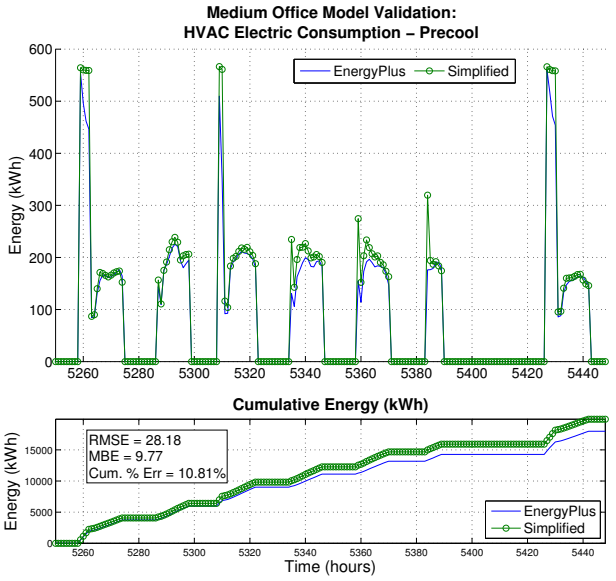

Figure 4.50: Medium office HVAC electric consumption for precooling scenario.

4.6.3 Large Office Model Comparison

The whole building validation was also repeated for the large office building, over a mild and hot validation week. The corresponding temperature, zone sensible load, and HVAC electric consumption are shown in Figures 4.51 – 4.56. The EnergyPlus validation data was generated from a detailed 15 zone representation. Overall, the simplified models performed well and were able to capture the desired dynamic response.

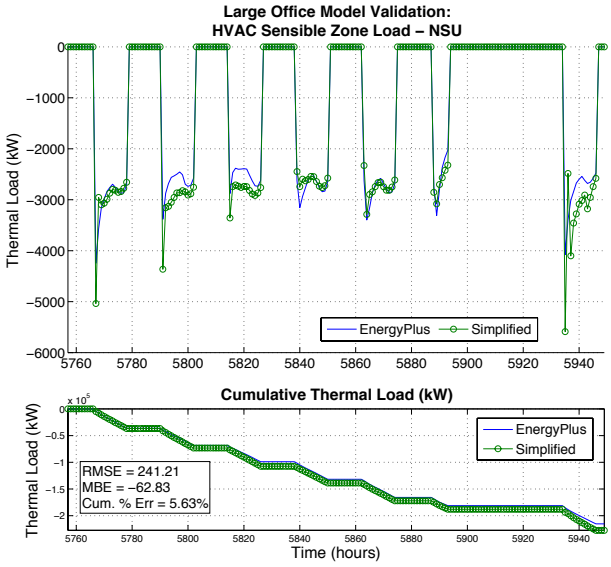

Figure 4.51: Large office sensible zone load for NSU scenario.

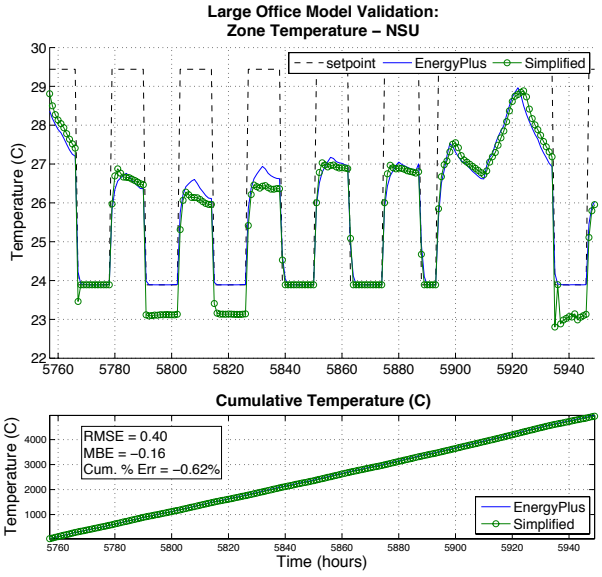

Figure 4.52: Large office zone mean air temperature for NSU scenario.

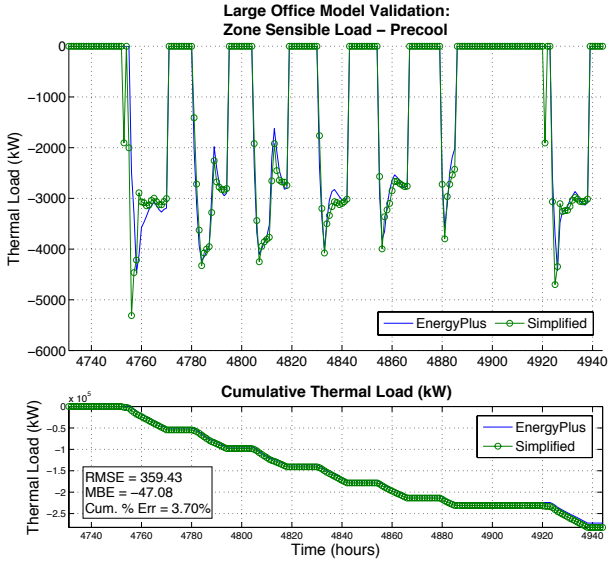

Figure 4.53: Large office sensible zone load for precooling scenario.

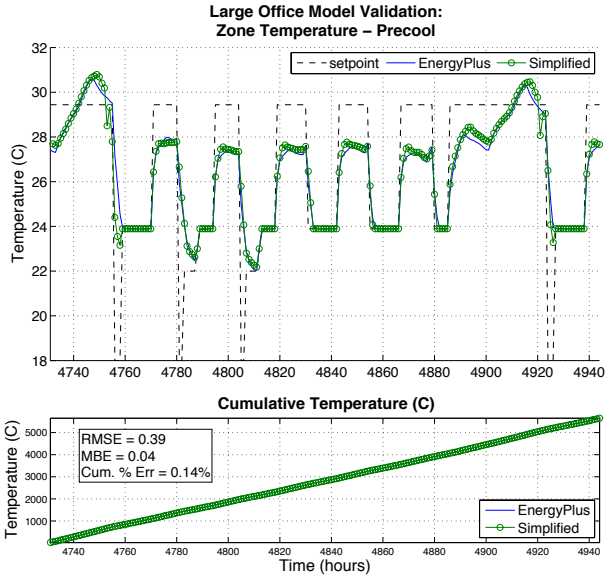

Figure 4.54: Large office zone mean air temperature for precooling scenario.

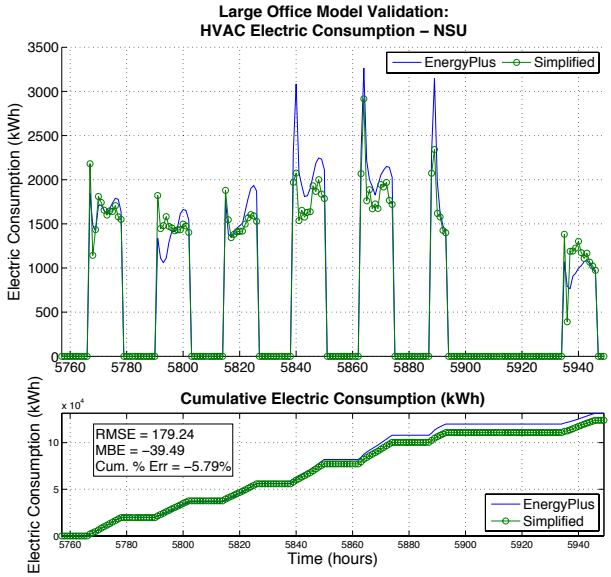

Figure 4.55: Large office HVAC electric consumption for NSU scenario.

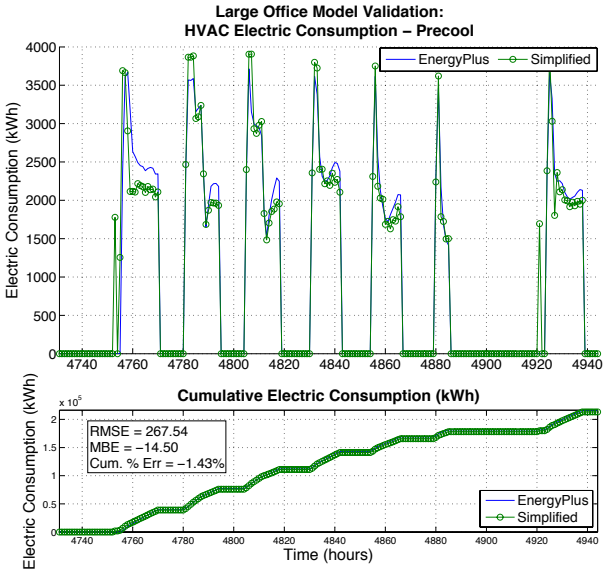

Figure 4.56: Large office HVAC electric consumption for precooling scenario.

4.6.4 Simulation with HVAC Models

The overall simulation scheme including HVAC can be described by the following "predictor-corrector" approach: 1) perform an ideal load calculation using Equation 4.3 to determine $Q_{zs,t}$ that meets the zone temperature setpoint , 2) simulate the HVAC system using the ideal load ($Q_{zs,t}$) as a requested system load, and 3) update zone temperature and load based on actual system performance and capabilities using Equation 4.4. This approach is repeated for each simulation time step allowing HVAC performance, control strategies, and capacity limits to affect zone conditions. The simulation also includes a simplified moisture balance to compute zone humidity levels as well.

For reference building characteristics are summarized for the three buildings below in Table 4.6.

Table 4.6: Building properties for retail, medium office, and large office models.

Property	Retail	Medium Office	Large Office	Units
Vintage	1980	2001	1980	year
Floors	1	6	32	#
Volume	13 984	59 028	256 808	m^3
Conditioned floor area	2294	14 240	76 659	m^2
U-value (no Film)	0.418	0.334	0.339	$W\,m^{-2}\,K^{-1}$
Internal Thermal Capacitance	450	3788	19 043	$MJ\,K^{-1}$
Internal Thermal Cap. per Floor area	196.2	266.0	248.4	$kJ\,K^{-1}\,m^{-2}$
Infiltration	3.93	2.1571	11.98	$m^3\,s^{-1}$
ACH	1.01	0.13	0.17	h^{-1}
Glazing Fraction	7	40	53	%
Glazing U-Factor	3.354	3.104	3.24	$W\,m^{-2}\,K^{-1}$
Glazing SHGC	0.385	0.306	0.498	fraction
Lighting Power Density	32.3	7.164	9.8	$W\,m^{-2}$
Equipment Power Density	5.23	4.5	4.63	$W\,m^{-2}$
Occupant Density	7.11	18.58	51.81	m^2/person

Chapter 5

Multi-Market Optimization

The development of an optimal multi-market scheduling methodology for commercial buildings presents some unique challenges not necessarily encountered by traditional generating plants or electric vehicles. First, a commercial building's primary responsibility is to serve the needs of its occupants by providing a healthy and productive environment. Satisfying this obligation requires any changes in building operation to maintain necessary thermal, visual, and indoor air quality requirements. These constraints are likely experienced during the majority of hours a building would provide ancillary services, since HVAC system operation and the most attractive frequency regulation prices typically correspond with occupied periods. Second, in the absence of active storage systems (e.g. battery banks or ice storage), a building's primary demand response (DR) asset is its passive thermal mass. The charging and discharging of a "thermal mass battery", and its instant effect on overall building electric consumption, is affected by ambient weather conditions, occupant behavior, interactions with HVAC systems, and the unique physical character of a building's thermal properties. Third, HVAC system capacity limits and efficiencies are a function of ambient and operating conditions and ultimately impact the ability to simultaneously achieve or maintain zone conditions and provide frequency regulation. Additional intricacy is created by equipment sequencing and staging logic. If such interactions are not anticipated, the response to FR signals could be periodically constrained and may trigger equipment startup and shutdown.

To accommodate such constraints and interactions, a whole-building model-based approach for estimating hourly commercial building regulating potential is proposed in this research. Ad-

ditionally, a MPC framework is applied to optimize building operation in consideration of both energy and regulation markets. The frequency regulation estimation methodology is described in Section 5.1, followed by a discussion of the multi-market optimization framework in Section 5.2. Sections 5.3, 5.4, and 5.6 discuss simulation study results, with a brief interruption in Section 5.5 to discuss the implications of the modeling methodology on the results.

5.1 Estimating Commercial Building FR Capability

To participate in frequency regulation markets it is necessary to quantify the continuous regulating capacity of a building during each bid period (often hourly). At a high-level the problem is similar to that posed in the electric vehicle literature in that it is desired to determine the potential of a building to increase or decrease its power draw with respect to a particular baseline (i.e. expected operating point). However, due to the complexity and challenges previously described, a neat linear formulation is precluded. As with generating plants, buildings have a nonlinear baseline since HVAC system operation changes intra-hour and hour-to-hour in response to varying electric prices, weather, and equipment staging. Similarly, potential regulation capacity varies throughout the day as a function, for example, of weather and how much HVAC equipment is operating in response to the weather at any moment. As an example, the response available from air handling units depends on how many are operating in any given hour and at what part-load ratio. Consequently, it is proposed to use a model perturbation approach to determine the relationship between building power response and changes in control input.

This concept is illustrated for FR via zone setpoint modulation in Figure 5.1. In the example, the model is assumed to be tracking a baseline zone temperature setpoint of 23 °C, and it is desired to determine the regulating capability for hour ending 1:00 PM. The model is used to evaluate the impact of zone setpoint changes from the baseline by simulating 0.1 K increments between a lower temperature limit of 22.22 °C (72 °F) and an upper limit of 23.89 °C (75 °F) as shown in the top window of Figure 5.1. Simulation results are compared with the original baseline demand profile to determine expected changes in whole building power. Increasing the temperature setpoint results

in a reduction in zone sensible load and HVAC power, providing regulation up. Decreasing the zone temperature setpoint increases HVAC power resulting in regulation down. The bottom window of Figure 5.1 shows the potential for approximately ±100 kW of power change by modulating the setpoint between the upper and lower limit. An advantage of using more complex models is that

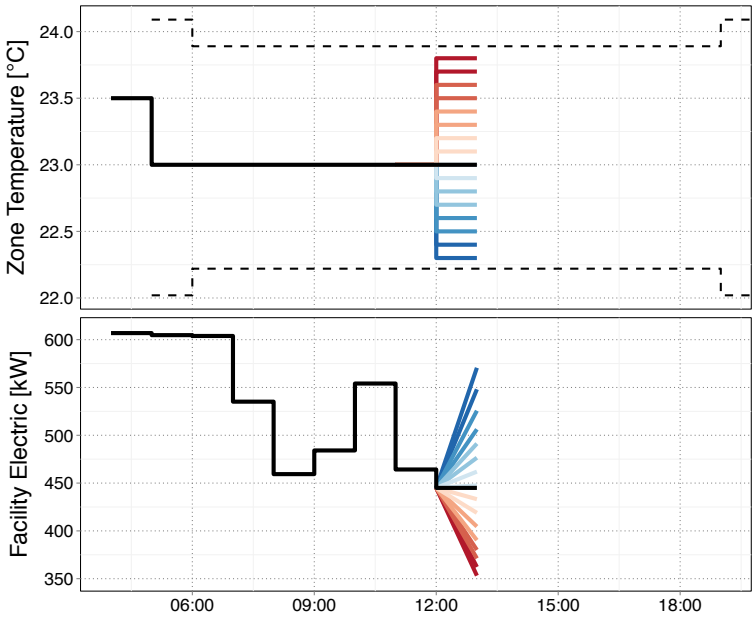

Figure 5.1: Zone setpoint perturbation example and the observed power response.

the power response can be grouped based on other simulation output variables. Depending on the HVAC system in use, it may be undesirable to allow the frequency regulation signal to cycle equipment on and off. In Figure 5.2 the power response is grouped based on the number of active cooling stages, and shows that limiting the modulation to a range of 22.6 °C to 23.4 °C would prevent turning stages on or off. Limiting the setpoint changes to this smaller range also reduces the regulating potential to approximately ±50 kW. The model perturbation analysis is repeated for each hour it is desired to determine the regulating capability. Figure 5.3 provides an example of performing the FR estimation for a period from hour ending 9 AM to 4 PM, assuming a constant baseline of 23 °C. The example highlights the fact that regulating potential may vary each hour depending on changes in HVAC equipment loading and efficiency.

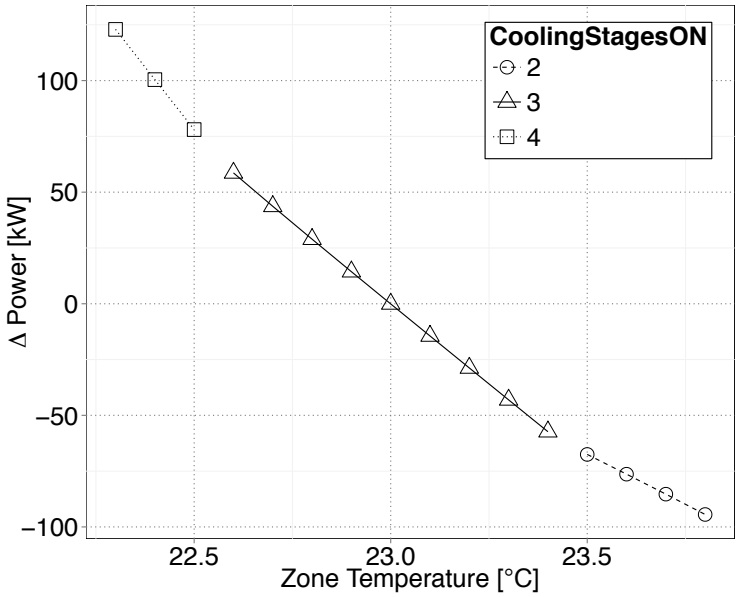

Figure 5.2: Setpoint perturbation power response grouped by active cooling stages.

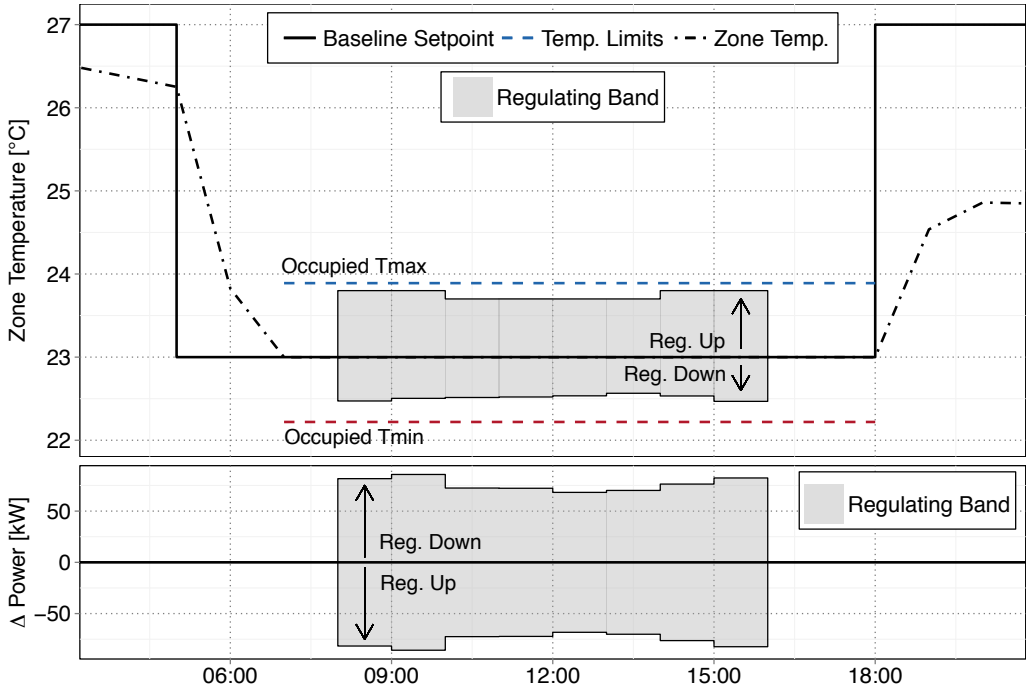

Figure 5.3: Setpoint perturbation performed for each hour from 9 AM to 4 PM.

The major assumption contained in this approach is that the effect of following the regulation signal, when integrated over each hour, does not cause significant deviation from the baseline setpoint strategy (i.e. the overall thermal strategy is maintained throughout the day). This assumption is ambitious when considering a single building since the regulation signal is random and can spend extended time in either direction. However, in the future context of building aggregates, the assumption seems more reasonable since the building aggregator may dispatch regulation duties over the course of all available hours in a manner that keeps each building near its preferred operating point and desired baseline strategy. Furthermore, the application of this work is intended more for the day-ahead planning and scheduling scenario. In buildings with large thermal inertia, load shifting strategies must often be planned and executed well in advance of peak or high priced hours. Estimating regulating capability and including in a day-ahead diurnal planning optimization allows buildings to be "prepared" for maximizing regulation, minimizing peak demand, maximizing economic demand response revenues, or an optimal blend of available opportunities.

5.2 Multi-market Optimization Framework

A model predictive control (MPC) framework was applied for optimizing building operation in the presence of real-time (or day-ahead) energy prices, peak demand charges, and frequency regulation revenue. The overall objective can be described as:

$$min\ J\left(\vec{x}\right) \qquad \text{s.t.: } \vec{x} \in [\vec{x}_{min}, \vec{x}_{max}]$$

where $\vec{x}$ is a control vector of setpoints in time, $\vec{x}_{min}$ is a vector of lower setpoint bounds, and $\vec{x}_{max}$ is a vector of upper setpoint bounds. The cost function $J\left(\vec{x}\right)$ is defined by Eq. 5.1,

$$J\left(\vec{x}\right) = E_{cost} + P_{demand} - R_{reg} \tag{5.1}$$

and includes terms for energy cost (E_{cost}), demand penalty (P_{demand}), and regulation revenue (R_{reg}). The cost of energy is computed by Equation 5.2,

$$E_{cost} = \sum_{t=1}^{t_{CH}} r_{DA}\left(t\right) E_{use}\left(t\right) \tag{5.2}$$

where $r_{DA}(t)$ is the day-ahead energy price for time interval t, $E_{use}(t)$ is the energy consumption over time interval t, and t_{CH} represents the final time interval in the cost horizon. Peak demand limitations are enforced by setting a target demand limit (TDL) and penalizing demand excursions above this limit. In this work, the linear demand penalty shown in Equation 5.3 was applied,

$$P_{demand} = \max\left[M\left(\max\left(ElecDemand_{peak}\right) - TDL\right), 0\right] \tag{5.3}$$

where TDL is the target demand limit and M is the slope of the penalty function. $ElecDemand_{peak}$ is a vector containing the average power for each 60-minute interval during the on-peak period. Demand charge calculations are ultimately utility specific, and may vary from this definition. Equation 5.3 should be modified to reflect location specific rules.

Regulation revenue (R_{reg}) is computed by summing the product of the potential power change (Δ_{power}) and regulation price (r_{reg}) for all time intervals (t) in the cost horizon (Equation 5.4).

$$R_{reg} = \sum_{t=1}^{t_{CH}} \Delta_{power}(t)\, r_{reg}(t) \tag{5.4}$$

It should be noted that Equation 5.4 is likely simplified from the actual calculation performed by regulation markets, and adjustments may be necessary based on ISO specific rules. In PJM, for example, a performance factor ranging from 0.25 to 1 is applied to scale the payment based on how well the regulation dispatch schedule was followed. Additionally, a "mileage" multiplier is included to appropriately compensate resources for following faster, dynamic regulation signals and providing more overall movement during the hour [83]. Therefore, Equation 5.4 could be thought of as the result of assuming a performance factor and mileage multiplier of unity under PJM rules.

This formulation considers the building as a "price taker," and assumes regulation prices are known in advance. Although perfect price forecasts were assumed in this work, numerous methods could be considered for price forecasting [34, 21, 112].

Figure 5.4 provides a graphical description of the multi-market optimization implementation. First, an initial baseline control vector $\vec{x}$ is generated by the optimizer. This initial control vector represents the baseline setpoint strategy or preferred operating point as discussed in Section 5.1. Second, the baseline strategy is simulated to determine the energy consumption and

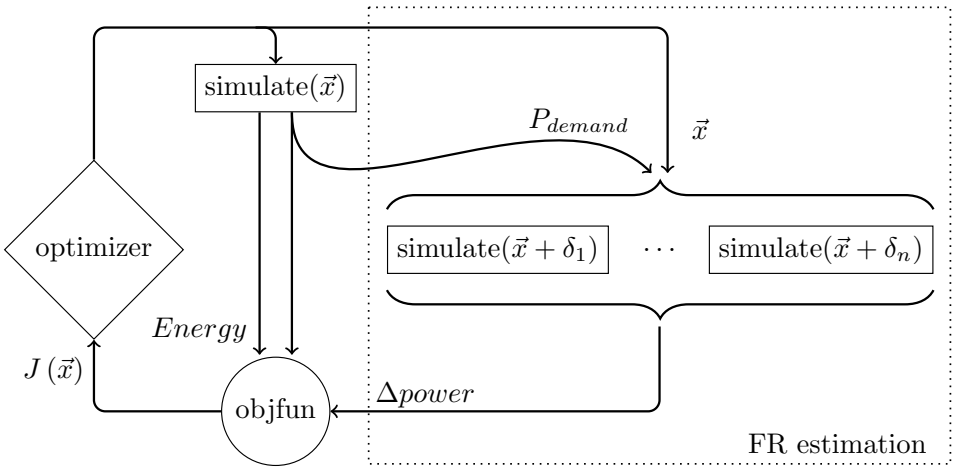

Figure 5.4: Overview of multi-market optimization with FR estimation.

demand penalty for that candidate control vector ($\vec{x}$). Third, the FR estimation is performed using the model perturbation approach to determine the potential variation around the baseline (i.e. Δ_{power}). Perturbed control vectors are denoted by the original baseline control vector $\vec{x}$ plus a setpoint change δ. The demand penalty (P_{demand}) from the baseline simulation is used in the FR estimation to ensure that the peak demand is not increased by providing regulation. Finally, the objective function of Equation 5.1 is computed, returning the total cost of the strategy back to the optimizer.

5.3 Medium Office Optimization

The multi-market optimization was applied the medium office building, developed in Chapter 4, for low and high target demand limit (TDL) cases. The low TDL case creates the scenario of a potentially peak setting day, while the high TDL case simulates a scenario in which peak demand is less critical. The building was assumed to be available for regulation from 9 AM to 5 PM, using the zone temperature setpoint as the modulation variable. Temperature bands were limited to a range that produces symmetric power response (i.e. equal regulation up and down) to satisfy ISO requirements. All FR simulation studies in this chapter were performed for the same June 24th, 2013, simulation period with ambient weather and solar conditions shown in Figure 5.5.

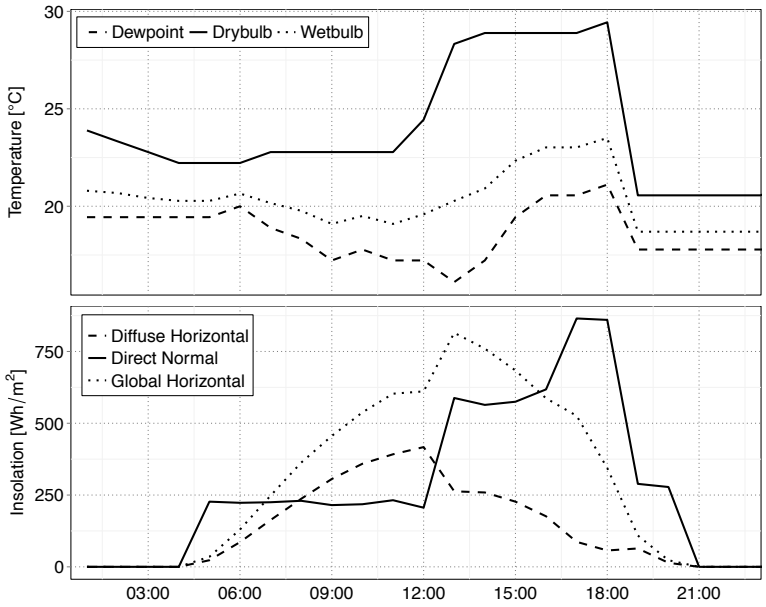

Figure 5.5: Ambient temperature and solar conditions for Chicago (June 24th, 2013).

Figure 5.6 shows the results for a low TDL of 325 kW (i.e. demand limiting). The night time setup (NSU) case represents typical building operation without optimization, while the OPT+FR case represents the result of the multi-market optimization process described in Section 5.2. In the OPT+FR case, the optimizer chose to start the HVAC system at hour ending 3 AM, precooling the zone to a temperature of 22.5 °C at hour ending 7 AM. The temperature is held near the middle and bottom of the occupied temperature range preceding the on-peak period to maintain stored thermal energy, and then stepped to the upper bound (Occupied T_{max}) during the on-peak period to release thermal energy back into the zone. With the OPT+FR baseline zone setpoint strategy, regulation was available for 5 of 9 potential hours ranging from ±8 kW to ±60 kW.

Overall, the primary objective of this scenario was to reduce on-peak demand since a significant $5.50 penalty is applied for each kW over the TDL. The cost function terms and energy use are summarized in Table 5.1, and show that expending an additional 364 kWh during early hours results in negligible increase in energy cost and significant demand savings. The minimal impact of regulation in this scenario is also shown, contributing only $12 of revenue. The temperature and demand profiles of the OPT case (i.e. zone setpoint optimization without regulation) are also

Table 5.1: Medium office multi-market optimization results.

	Low TDL			High TDL		
	NSU	OPT	OPT+FR	NSU	OPT	OPT+FR
Energy Cost [$]	246.51	247.65	250.93	246.51	240.82	264.15
Energy Use [kWh]	5175	5596	5539	5175	5257	5534
Demand Penalty [$]	320.27	103.66	103.99	0	0	0
Reg. Revenue [$]	0	0	-11.87	0	0	-56.39
$J(\vec{x})$ [$]	566.78	351.31	343.05	246.51	240.82	207.76
% Difference	0	-38.02	-39.47	0	-2.31	-15.72

plotted for reference and show a similar strategy when regulation opportunities are excluded.

The results for the high TDL scenario are shown in Figure 5.7. Looking at the OPT temperature and demand profiles shows an earlier start-up with slight precooling before 6 AM, followed by near NSU operation for the remaining hours. Table 5.1 shows that expending an extra 82 kWh earlier in the day results in slight overall energy cost savings ($5.69) by reducing consumption during high-price hours. No regulation would be possible in the OPT case since the temperatures remain near the upper bound during hours available for regulation. In the OPT+FR case, the optimizer chose to follow the NSU strategy until hour ending 8 AM, after which the building is available to provide regulation. During regulation hours the baseline zone setpoint was kept in the middle of the acceptable temperature range in order to create room for providing regulation. Keeping the zone setpoint 1 K lower during regulating hours uses 360 kWh more than the NSU case, however the regulation revenue more than offsets the additional energy cost resulting in an overall savings of 15.7%. This additional energy expenditure to create flexibility in the zone temperature could be considered an opportunity cost of providing FR. On average, an estimated ±85 kW of regulation were available, with a maximum of ±100 kW available during hour ending 10 AM.

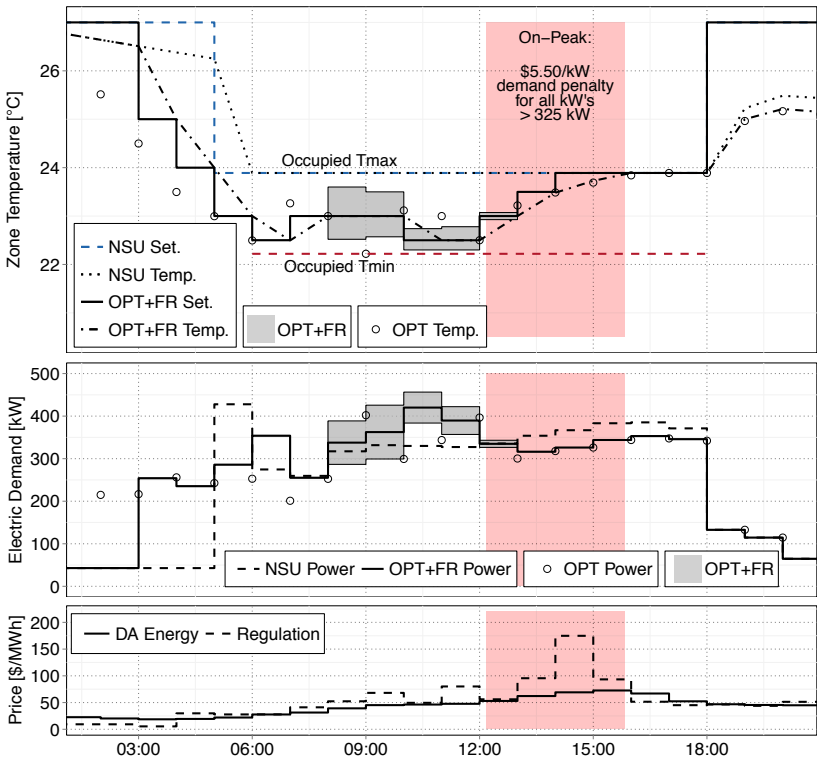

Figure 5.6: Medium office zone setpoint optimization with FR and low TDL.

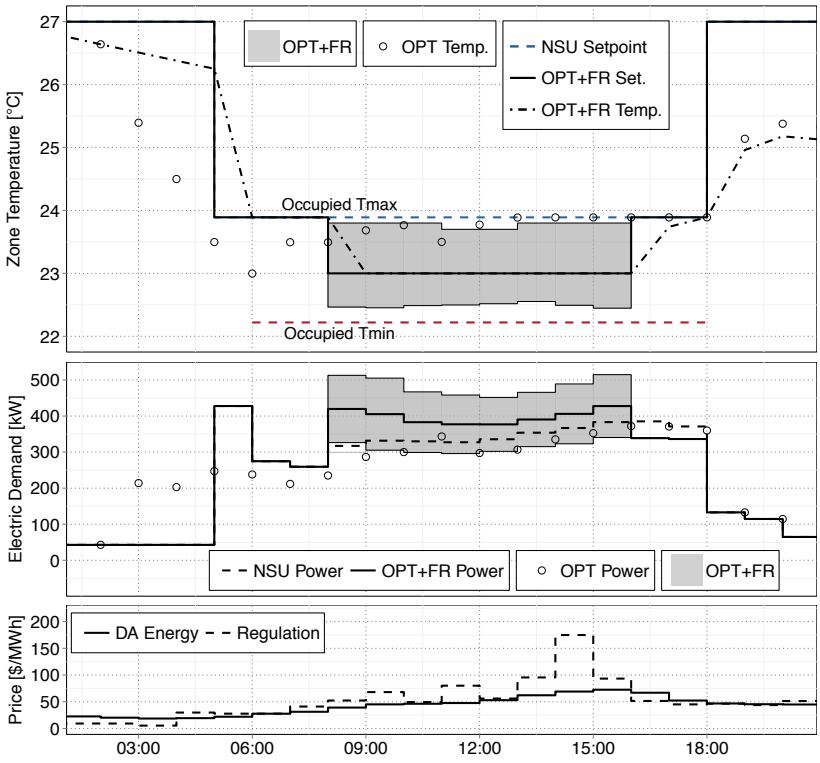

Figure 5.7: Medium office zone setpoint optimization with FR and high TDL.

5.4 Large Office Optimization with Steady Response FR Estimation

The multi-market optimization was also applied to low TDL and high TDL cases for the large office building, developed in Chapter 4, served by a chilled water VAV system. For these cases the optimizer was given control over the zone temperature setpoint to take advantage of passive thermal mass, as well as the chilled water (CHW) temperature setpoint to alter chiller plant efficiencies and loading (i.e. the optimizer chooses both zone and CHW temperature setpoints). This building was assumed to be available to provide FR from 9 AM to 5 PM, using the CHW setpoint as the modulation variable. In this example, the chilled water supply is typically maintained at a constant 5.56 °C. It was assumed that an operating range of 3.89 °C to 10.00 °C could be explored for FR without detrimental effects to equipment. Since increasing the chilled water temperature can result in the inability to maintain supply air temperature (and consequently zone temperatures) a large penalty is added to solutions that result in degradation of zone conditions. A non-cycling condition was included in the FR estimation so that providing regulation does not turn chillers on or off. The requirement of symmetric regulation was also enforced.

The top panel of Figure 5.8 shows that a zone precooling strategy is adopted that maintains the lower setpoint bound until the on-peak period. During the demand-limiting period the setpoint is raised to the upper bound to allow discharging of the thermal mass and reduction of on-peak demand and energy consumption during high-priced hours. The second panel from the top of Figure 5.8 shows the chosen baseline chilled water temperature setpoints with regulating bands. Unlike the zone setpoint FR example, the CHW FR NSU case has a constant baseline between the temperature limits and is able to participate in regulation during 8 of 9 available hours. Regulation is not possible during hour ending 4 PM for the NSU case since the peak demand is set during this hour. The NSU regulation capability ranges from ±30 kW to ±70 kW with an average of ±50 kW, as shown in the third panel from the top of Figure 5.8. The OPT+FR case is able to participate in the two hours immediately preceding the on-peak period, with other hours unavailable due to the baseline CHW setpoint being at the upper or lower bound and the requirement to provide symmetric

regulation. Although, the chosen zone temperature trajectory is virtually identical between the OPT and OPT+FR cases, the OPT case baseline would not be able to provide regulation due to temperatures always being at a lower or upper limit during regulating hours. Overall, the solution is primarily seeking to reduce on-peak demand and the optimizer determines it is beneficial to forego several hours of regulation revenue to achieve these demand savings.

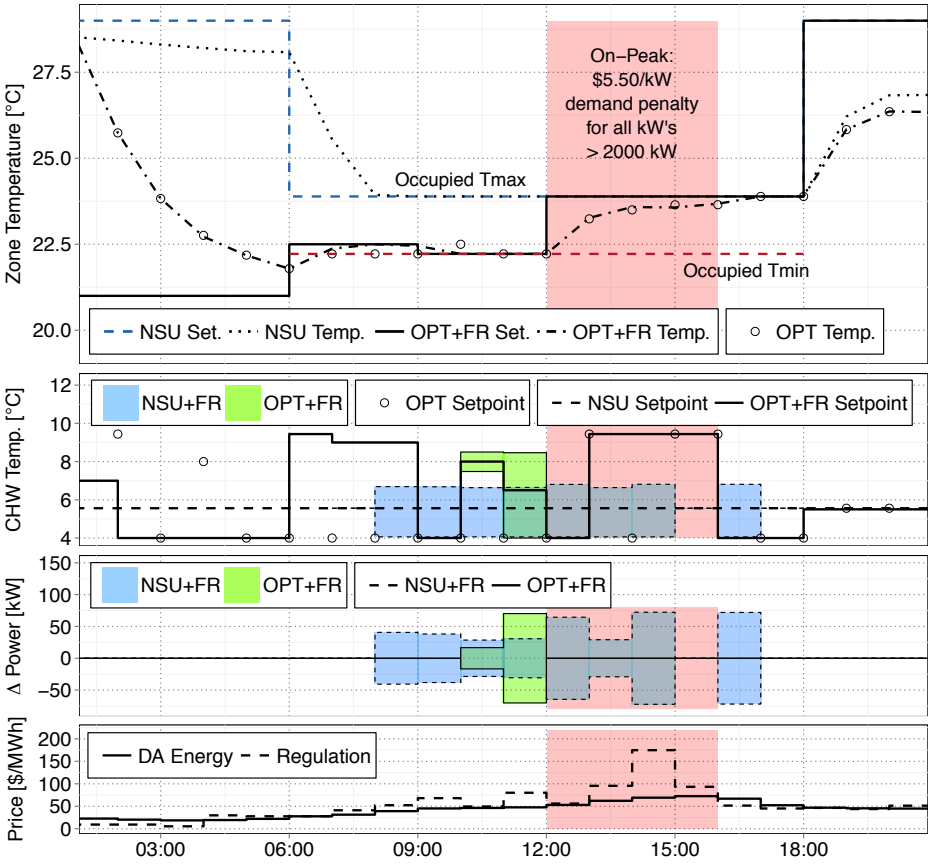

Figure 5.8: Large office optimization with low TDL and steady response FR estimation.

The results for the large office scenario with high TDL are shown in Figure 5.9. The top panel shows that the OPT+FR case chose a slightly early startup and followed NSU operation during occupied hours. The OPT+FR CHW setpoint tended to stick near the upper bound with a few hours near the middle of the range during regulating hours. The early start-up and higher average CHW temperatures result in overall energy and cost savings, presumably due to economizing and

more efficient loading and operation of the chiller plant. The OPT+FR case found it beneficial to participate in regulation for 5 of 9 hours with an average capability of ±54 kW. The NSU results are similar to the low TDL case with the exception of hour ending 4 PM being available for FR since the demand limit is no longer a factor.

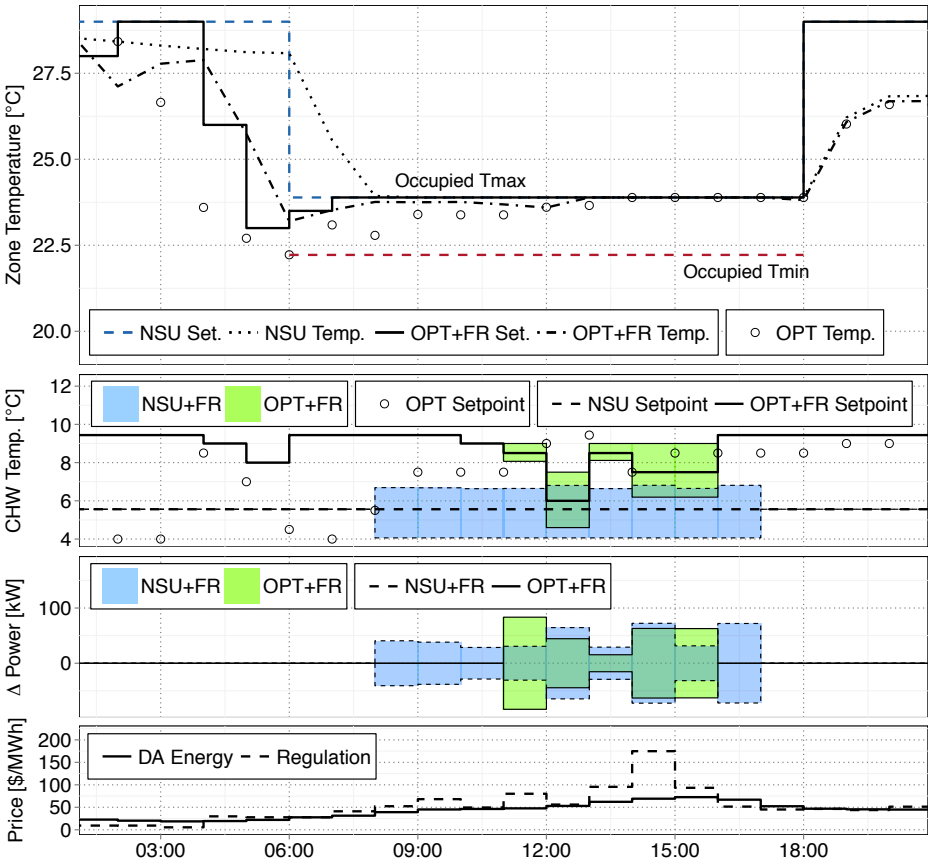

Figure 5.9: Large office optimization with high TDL and steady response FR estimation.

Overall, the estimated FR potential for this example was similar to the medium office building one-fifth the size, and an average regulating potential of ±50 kW seems somewhat unimpressive for an 8600 kW cooling plant with a rated power draw near 2 MW. To better understand these results and estimates, a further discussion on the modeling approach is provided in the following section.

5.5 An Interlude On Modeling

The modeling approach taken in this work makes use of a reduced-order thermal zone model and quasi steady-state HVAC models simulated at hourly time steps. The result is that hourly simulation output reflects the conditions reached by the end of the hourly time step. It is necessary to further consider the consequence of using such hourly output for FR estimation purposes.

In the context of FR via zone temperature setpoint modulation (e.g. Section 5.3), decreasing the zone cooling setpoint increases the zone cooling load and results in increased HVAC power. Comparing hourly outputs in this scenario results in FR estimates that include full HVAC system response due to changes in zone load, since system transients typically settle at timescales less than one hour. In this case, FR estimates represent a maximum power change that is not necessarily instantaneously achievable at the two second FR timescale. In reality some delay would be associated with the full HVAC system response and it may be necessary to apply a derate factor to adjust maximum estimates based on observed building response to regulation signals.

However, when estimating FR capability via CHW setpoint modulation (e.g. large office results of Section 5.4), comparing the hourly simulation output may significantly underestimate the regulating potential. The difference, with respect to the zone setpoint case, is that changes in CHW setpoint ultimately do not affect the overall load placed on the cooling coil. It is expected that changes in CHW setpoint would initially excite a power response due to the transient loading or unloading of the chiller evaporator. This power response would then degrade as CHW return temperature is impacted and the temperature differential across the evaporator moves back toward steady-state conditions. At the new steady-state, some differences in HVAC power may still exist since chiller and cooling coil performance are a function of CHW temperature. CHW supply temperature impacts the sensible and latent capacity of the cooling coil and can somewhat influence the total plant load since, in this case, the coil is controlled based on sensible performance only (i.e. supply air temperature). In this scenario, using near steady-state hourly outputs to estimate FR would only reflect power changes due to differences in chiller and cooling coil performance,

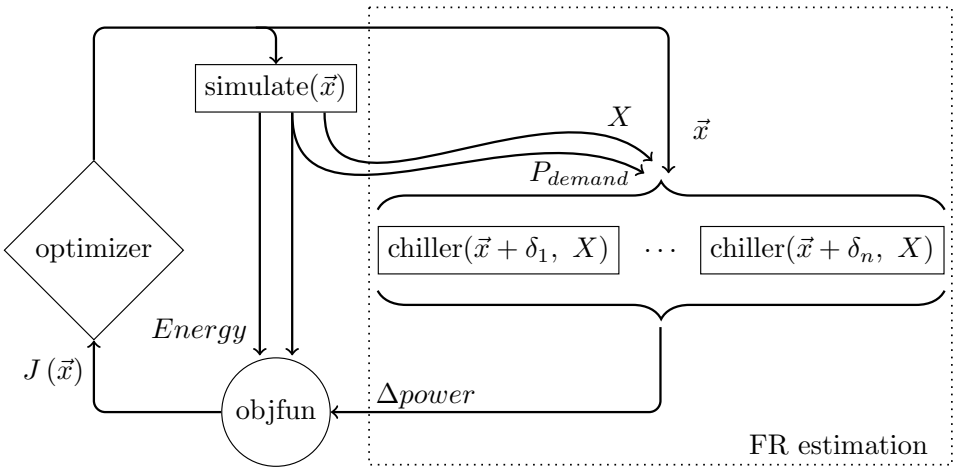

Figure 5.10: Overview of multi-market optimization with isolated chiller model FR estimation.

rather than evaporator load.

In order to estimate potential chiller response due to transient loading, the CHW FR estimation methodology was modified to use an isolated chiller model, rather than simulating the entire plant. The chiller model used throughout this work has the functional form:

$$P_{chiller} = f\left(T_{chw,sp},\ T_{chw,r},\ \dot{m}_{chw},\ T_{cdw,in},\ \dot{m}_{cdw}\right)$$

where $T_{chw,sp}$ is the CHW supply temperature setpoint, $T_{chw,r}$ is the CHW return temperature, $\dot{m}_{chw}$ is the CHW mass flow rate, $T_{cdw,in}$ is the condenser water inlet temperature, and $\dot{m}_{cdw}$ is the condenser water mass flow rate. Simulating only the chiller model allows all flow rates and inlet temperatures to be held constant while changing the CHW setpoint. This results in estimates that include changes in chiller performance as well as overall evaporator load.

These changes within the optimization are graphically summarized in Figure 5.10. First, the optimizer generates a candidate baseline setpoint vector $\vec{x}$ and simulates to determine the baseline performance. Second, the baseline control vector $\vec{x}$, demand penalty P_{demand}, and chiller states X (i.e. $T_{chw,r}$, $\dot{m}_{chw}$, $T_{cdw,in}$, and $\dot{m}_{cdw}$) are passed along for use in the FR estimation. Third, the chiller model is simulated changing only the chilled water supply temperature, resulting in the available power change (Δ_{power}). Finally, the overall objective function is computed returning total cost $J\,(\vec{x})$ to the optimizer.

The estimates produced from this method may be closer to the maximum regulating potential and are not necessarily achievable instantaneously. As with the zone setpoint method, a degradation factor may be necessary to scale estimates based on observed building response. Herein, the FR estimation using the isolated chiller model is referred to as the "fast response" method, and the large office FR estimation using the whole building hourly results is referred to as the "steady response" method.

5.6 Large Office Optimization with Fast Response FR Estimation Method

Results for the large office low TDL multi-market optimization using the fast response FR estimation method are shown in Figure 5.11. Overall, the results are similar to the steady response method with the exception that regulation estimates are much higher. On average, the fast response method estimated ±450 kW of regulation compared to ±50 kW by the steady response method for the OPT+FR case. Although the NSU+FR operation could have generated over $200 of revenue for the day, greater priority is still given to demand reduction due to the steep penalty. Energy use and cost function terms for the large office low TDL cases are summarized in Table 5.2.

In Figure 5.12, the high TDL optimization results using the fast response method show similar OPT and OPT+FR solutions for the zone temperature setpoint when compared to the previous results in Figure 5.9. Comparing the CHW setpoints shows that when using the fast response method the CHW temperature was kept near the middle of the temperature range in order to provide more regulation. Since the fast response method estimates greater regulating potential, the revenue generated is greater than any savings generated from keeping a higher CHW temperature. The OPT+FR was able to provide regulation during 8 of 9 hours, with a maximum of capability of ±1 MW at 3 PM. The average capability was ±610 kW, generating nearly $465 in revenue for the day. The energy use and cost function terms for the high TDL cases are summarized in Table 5.3. It is further noted that although several scenarios resulted in higher average CHW temperatures, zone humidity levels remained reasonable. The maximum zone humidity ratio experienced during occupied hours in any of the large office results was 9.6 g kg^{-1}, with an average of 8.5 g kg^{-1}.

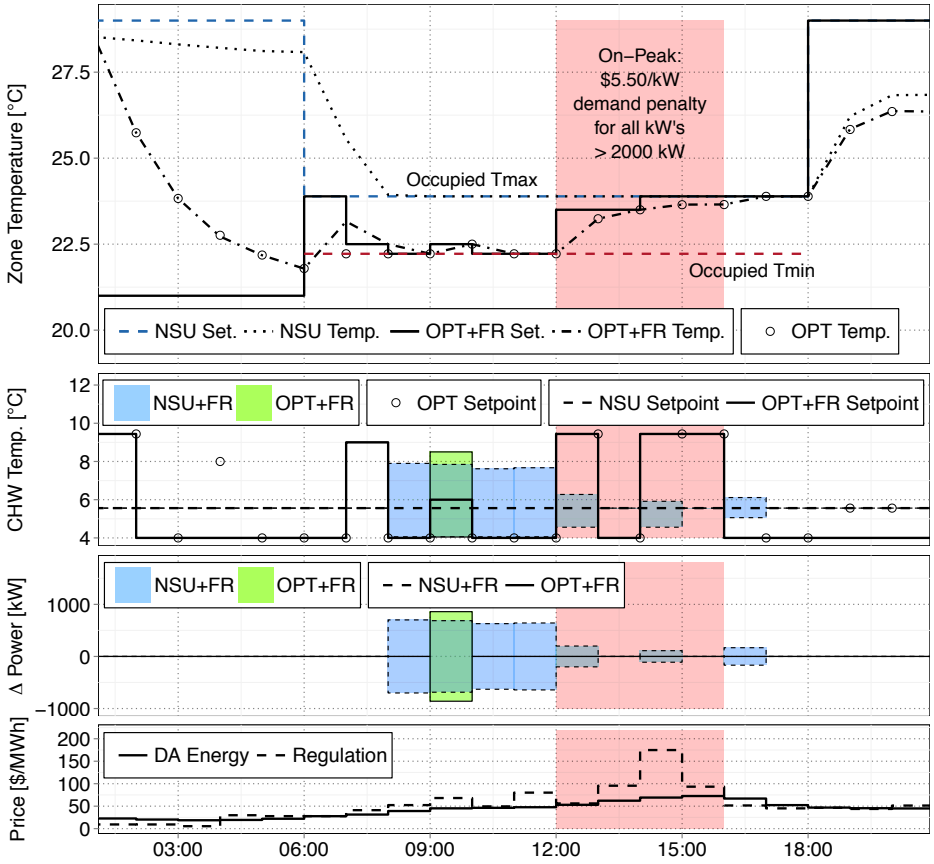

Figure 5.11: Large office optimization with low TDL and fast response FR estimation.

Table 5.2: Large office low TDL results.

			Steady Response Method		Fast Response Method	
	NSU	OPT	NSU+FR	OPT+FR	NSU+FR	OPT+FR
Energy Cost [$]	2061.36	2070.83	2061.36	2031.82	2061.36	2063.54
Energy Use [kWh]	42 861	51 908	42 861	51 025	42 861	51 548
Demand Penalty [$]	5573.92	1585.69	5573.92	1285.99	5573.92	1652.02
Reg Revenue [$]	0	0	-31.36	-6.45	-205.41	-58.58
$J(\vec{x})$ [$]	7635.28	3656.53	7603.92	3311.35	7429.87	3656.99
% Difference	0	-52.11	-0.41	-56.63	-2.69	-52.10

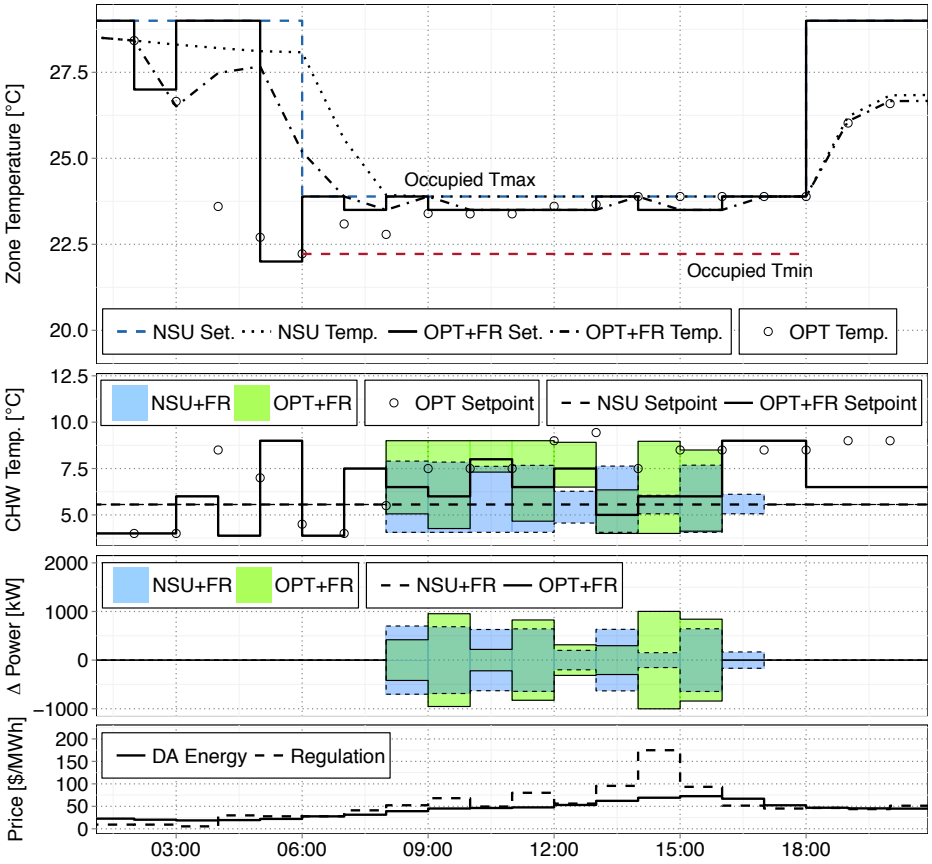

Figure 5.12: Large office optimization with high TDL and fast response FR estimation.

Table 5.3: Large office high TDL results.

			Steady Response Method		Fast Response Method	
	NSU	OPT	NSU+FR	OPT+FR	NSU+FR	OPT+FR
Energy Cost [$]	2061.36	1836.78	2061.36	1855.31	2061.36	2023.74
Energy Use [kWh]	42 861	43 474	42 861	40 829	42 861	43 944
Demand Penalty [$]	0	0	0	0	0	0
Reg Revenue [$]	0	0	-34.30	-27.51	-333.51	-463.73
$J(\vec{x})$ [$]	2061.36	1836.78	2027.06	1827.80	1727.85	1560.01
% Difference	0	-10.89	-1.66	-11.33	-16.18	-24.32

Chapter 6

Portfolio Optimization Development and Testing

6.1 Portfolio Optimization Environment

In pursuit of exploring the existence of synergy through optimal control of building portfolios, it was necessary to extend the single building model predictive control environment, originally developed by Corbin et al. [24], to accommodate multiple buildings. Optimizing a portfolio of N buildings can be considered a generalization of the single building problem, and involves managing initialization, optimization, and execution tasks across multiple buildings. A graphical representation of the portfolio optimization is provided in Figure 6.1, and further discussion is provided in the following subsections.

6.1.1 Initialization

The first phase of the model predictive controller performs an initialization of the simulation and optimization parameters. The generalization to multiple buildings was created through repetition of lower level tasks for all building. This process is illustrated through the `for` loop pseudocode shown in the top section of Figure 6.1.

The first task involves loading building specific parameters to update high-level defaults. It is convenient to categorize parameters as being global in the sense that they apply to all buildings equivalently, or as local parameters that describe unique instructions for each building. Parameters specifying optimization convergence tolerances and iteration limits, for example, can be considered global and need not be uniquely defined for each building. It was also assumed in this work that

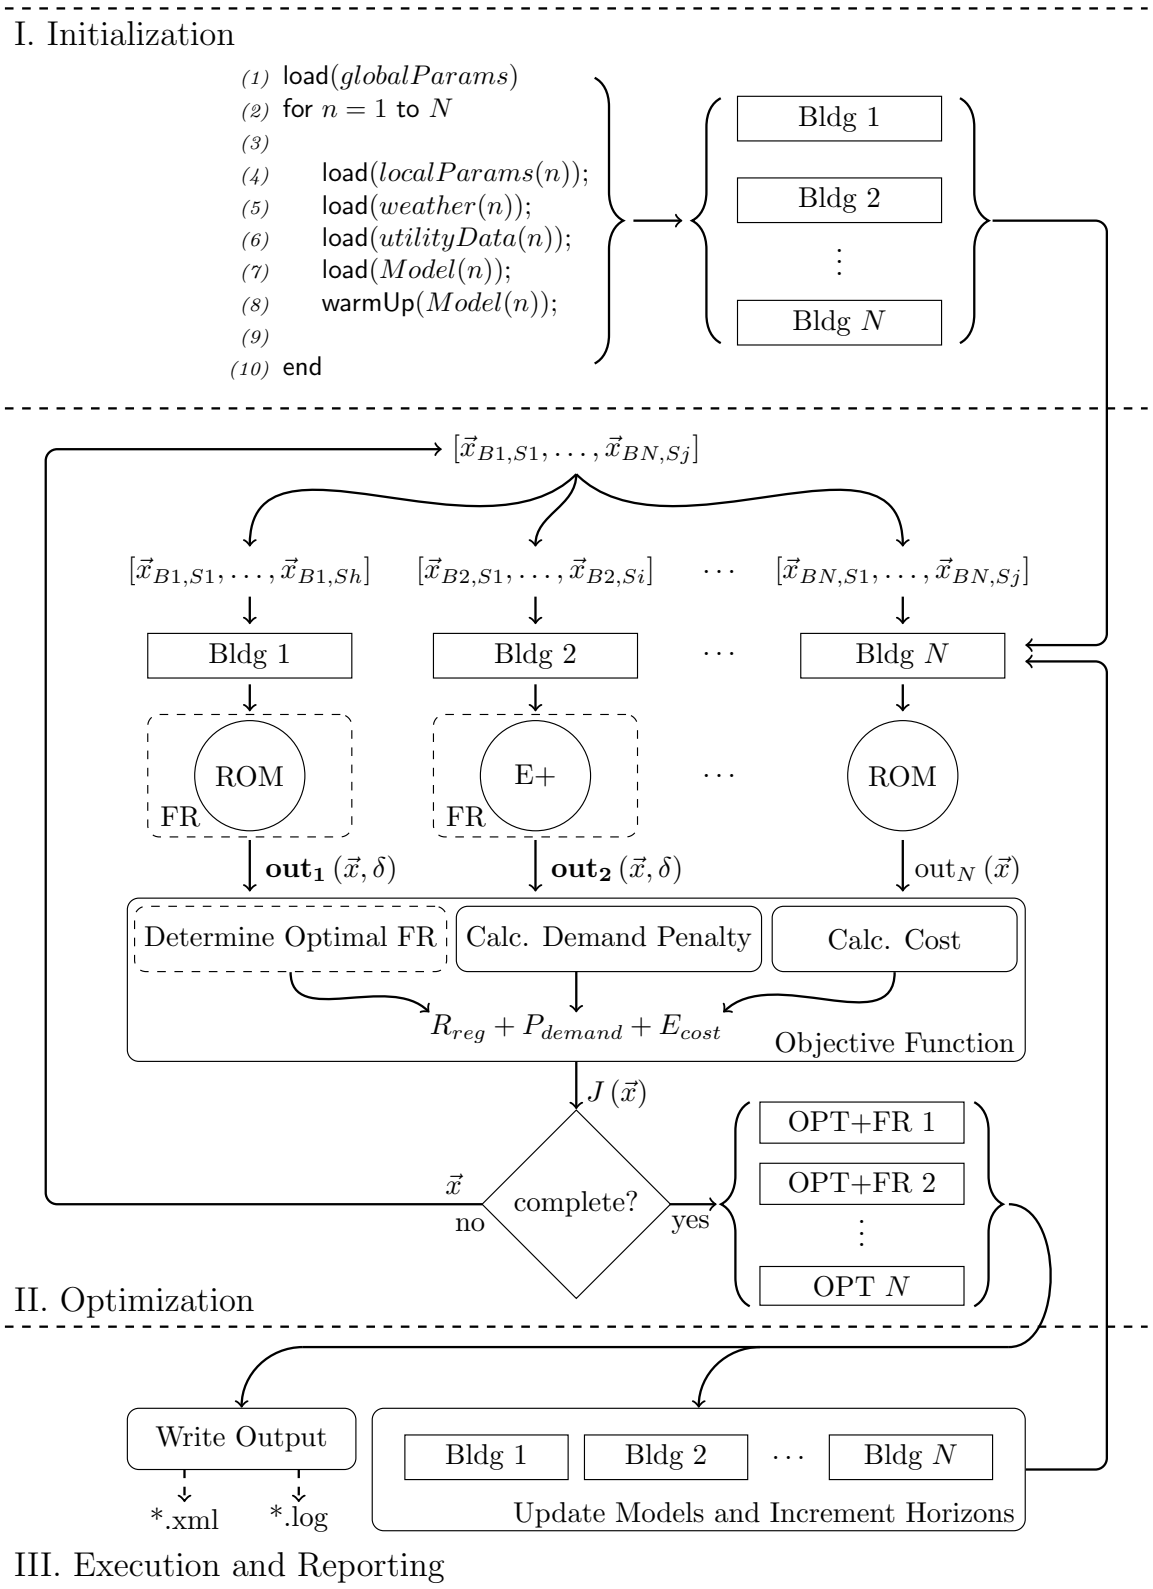

Figure 6.1: Diagram showing multi-market portfolio optimization framework.

all optimized buildings follow the same MPC horizon structure, having equal length planning and execution windows. In order to consider a wide variety of cases, it is necessary to allow optimization variables, discretization, and variable bounds to be uniquely defined for each building. Additional examples of building-specific parameters are the simulation engine, comfort penalty, and availability for frequency regulation.

In the second task outlined under the initialization phase, weather forecasts and history are loaded for each building. Weather has been considered building-specific in order to provide a more generalized framework. When optimizing a portfolio of co-located buildings it seems likely that weather forecasts would vary minimally between buildings. However, measured historical data may be available if buildings are equipped with weather stations, which would capture local variations between buildings. Historical weather data is important for establishing an accurate thermal history within the model, keeping the model as near to reality as possible. Although not considered in this work, it may be desired to optimize dispersed buildings which experience entirely different weather patterns as well.

In conjunction with weather, utility data is loaded for each building in the third task of the initialization phase. Since the buildings operate independently, and cannot share thermal energy, their energy consumption is independent. Thus, energy expense can be computed using different energy price structures for each building within the portfolio optimization. Ancillary service prices are often established for an entire balancing area or region and may not be unique for co-located buildings. However, such prices could also be defined uniquely for each building as well.

The final tasks of initialization involve loading building energy model descriptions specific to the desired simulation engine, and pre-establishing model thermal history if appropriate. The reduced order models used in this work allow building thermal states to be prescribed and propagated through optimization iterations, which allows the building "warm up" to be performed only once during initialization. However, when using EnergyPlus as the simulation engine the thermal history must be re-established at each simulator call. Further information regarding establishment of thermal history within EnergyPlus is described by Corbin et al. in [24]. The initialization phase

produces an organized container of data and parameters for each building to be used and modified throughout the MPC operation. In Figure 6.1, this is represented by the black rectangles labeled "Bldg 1," "Bldg 2," and "Bldg N."

6.1.2 Optimization

The MPC optimization was performed using a variant of the metaheuristic particle swarm (PSO) algorithm, as implemented by Corbin et al. [24]. The algorithm generates a candidate control vector and expects the cost of that control strategy to be returned as an input. It is indifferent to how the cost is calculated, thus, the extension from single building to portfolio optimization did not require any modification of the original optimization routine.

The control vector $\vec{x}$ provided by the optimizer contains the decision variables for all buildings in the portfolio. This can be denoted as $[\vec{x}_{B1,S1}, \ldots, \vec{x}_{BN,Sj}]$, where subscript $B1, S1$ represents the first control schedule for building one, and BN, Sj represents the j-th control schedule for building N. As shown in the optimization phase of Figure 6.1, this initial control vector must be split into separate decision vectors for each building and populated into the corresponding model. Each building may desire to optimize a different number of schedules which is denoted through the use of separate schedule indices h, i, and j in the figure. After decision variables are divided among buildings, the building model and parameters are passed to the simulation engine(s) to evaluate the control vectors. The environment currently only allows for a combination of reduced order and EnergyPlus models to be used within the same portfolio optimization, however, similar extensions could be made to include other simulators as well.

For a particular building it may also be desired to include the estimation of the frequency regulation potential, as denoted in Figure 6.1 by the dashed rectangles labeled "FR." For further detail on the methods used in estimating FR capability see Chapter 5. Simulation results are then passed to the objective function for use in computing the total cost of the control vector. The simulation outputs for Bldg 1 and Bldg 2 are in bold since they are a function of both the initial control vector $\vec{x}$ and FR perturbations δ. These results contain a set of separate simulation outputs

for each value of δ, which is necessary for determining the optimal amount of FR. The objective function can also be computed with the inclusion of building multipliers that scale the output of a single simulation. This allows numerous portfolio combinations to be evaluated without increasing the optimization complexity.

Within the objective function, the portfolio energy cost is computed by summing the individual energy expense of all buildings as shown in Equation 6.1,

$$E_{cost} = \sum_{n=1}^{N} \sum_{t=1}^{t_{CH}} r_{DA,n}(t)\, E_{use,n}(t) \tag{6.1}$$

where $r_{DA,n}(t)$ is the day-ahead energy price at time t for the n-th building, $E_{use,n}(t)$ is the energy consumption over time interval t for the n-th building, and t_{CH} represents the final time interval in the cost horizon.

For the portfolio optimizations, an enhancement was made to the FR revenue calculation methods presented in Chapter 5. In the single building FR estimation the assumption was made that it was undesirable to increase peak demand in order to provide frequency regulation. Since the demand penalty of \$5.50/kW was orders of magnitude larger than the regulation price, this assumption was reasonable. Any FR perturbations that exceeded the target demand limit were discarded, leaving only the maximum FR capability below the demand threshold. However, if the demand penalty were of similar magnitude as ancillary service prices, it seems that an economic decision could be made to trade higher demand penalties for increased revenue. Furthermore, increasing the peak demand in one hour raises the demand limit for the entire on-peak period, potentially allowing more revenue to be generated during all on-peak hours.

In the context of building portfolios the situation becomes even more complex. Multiple buildings are contributing to an aggregate demand and an increase in the peak demand would allow all buildings to potentially generate more revenue during all on-peak hours. In the event it is not economical to increase peak demand and provide maximum regulation it must be determined how much frequency regulation each building should contribute. Due to capacity limitations and operational constraints that may not occur coincident in time among buildings in the portfolio,

there may be opportune times for each building to provide specific amounts of regulation. In order to create a broad framework that can appropriately weigh the economic value of regulation and demand penalties under any given price scenario, it is necessary to define a second optimization problem. The sub-optimization, in essence, seeks to determine the optimal amount of frequency regulation considering the cost of peak demand excursions.

In this work the FR optimization was approached using "brute force" by computing the sum of demand penalties (P_{demand}) and regulation revenue (R_{reg}) for all combinations of FR perturbations. Performing the FR perturbation analysis for each building results in a separate set of output variables for each perturbation δ. The electric demand profiles for all buildings and all perturbations can be enumerated exhaustively to create a single matrix that contains portfolio demand profiles for all potential FR scenarios, shown in Equation 6.2. The columns of the matrix represent simulation time steps while each row is a realization of a single perturbation combination. The example matrix shown is for two buildings where δ_{21} represents the building 2 control vector with perturbation 1, and indices f and g represent the final perturbations for $B1$ and $B2$, respectively.

$$\mathbf{ElecDemand} = \begin{bmatrix} D\left(1, \delta_{11}, \delta_{21}\right) & D\left(2, \delta_{11}, \delta_{21}\right) & \dots & D\left(t, \delta_{11}, \delta_{21}\right) \\ D\left(1, \delta_{11}, \delta_{21}\right) & D\left(2, \delta_{11}, \delta_{21}\right) & \dots & D\left(t, \delta_{11}, \delta_{22}\right) \\ \vdots & \vdots & \ddots & \vdots \\ D\left(1, \delta_{1f}, \delta_{2g}\right) & D\left(2, \delta_{1f}, \delta_{2g}\right) & \dots & D\left(t, \delta_{1f}, \delta_{2g}\right) \end{bmatrix} \tag{6.2}$$

Demand penalties can be computed for all portfolio demand profiles as in Equation 6.3,

$$P_{demand} = \textsc{max}\left[M\left(\textsc{max}\left(\mathbf{ElecDemand}\left(:, t_{peak}\right)\right) - TDL_p\right), 0\left(:, 1\right)\right] \tag{6.3}$$

where t_{peak} represents a time index spanning all on-peak hours, TDL_p represents the portfolio target demand limit, M represents the slope of the demand penalty in units of \$/kW, and $0\left(:, 1\right)$ represents a column vector of zeros. Small capital letters are used in the "MAX" function to denote

a matrix function that operates on each row. Equation 6.3 results in a column vector of demand penalties (P_{demand}) for all portfolio demand profiles.

Matrix operations can similarly be used in computing the regulation revenue for each combination of perturbations. The baseline portfolio demand (i.e. without perturbation) can be subtracted from the absolute demand profiles to determine the power change available for regulation, denoted as **ΔElecDemand**. The regulation revenue (R_{reg}) can than be computed by

$$R_{reg} = \text{SUM}\left(\mathbf{\Delta ElecDemand} \circ \mathbf{r_{reg}}\right) \tag{6.4}$$

where $\mathbf{r_{reg}}$ represents a matrix of regulation prices with size equal to **ΔElecDemand**, and SUM represents a sum over matrix rows. The "∘" operator is used to denote element wise matrix multiplication. The overall objective of Equation 6.4 is to compute the regulation revenue for all perturbation combinations listed in Equation 6.2. The demand penalty vector and regulation revenue vector can then be added to determine the combination that leads to the lowest cost. The overall cost of the control vector $J(\vec{x})$ is computed by adding the energy cost, demand penalty, and regulation revenue as in Eq. 6.5.

$$J(\vec{x}) = E_{cost} + P_{demand} - R_{reg} \tag{6.5}$$

The optimization continues until an exit criterion is met, at which point the optimal results for each building are passed along to the execution and reporting phase.

For the portfolio configurations and discretization used in this work, **ElecDemand** is typically around 300 000 rows, by 24 colums (i.e. hourly). Given an adequate amount of computer memory, basic matrix operations with variables of such dimension can be performed relatively quickly. The majority of additional compute time required to perform an FR analysis is mainly attributable to simulation of control vector perturbations. However, since the perturbations are independent they can be executed in parallel to significantly reduce execution times. The exhaustive enumeration approach was adopted in this work since it could be performed in a reasonable amount of time and guarantees an optimal solution. When scaling to higher dimensional problems,

with more buildings, this approach will likely be infeasible and optimization techniques should be considered.

6.1.3 Execution and Reporting

During the execution and reporting phase, the optimal solutions for each building are be written to various output files. Solutions are executed in the MPC controller by adding the optimal control vectors and thermal states to the building model histories. This work was performed entirely within a simulation testbed, thus the solutions were essentially executed by implementing them in the same model used for optimization (i.e. no model mismatch). The MPC horizons are then incremented in time to begin optimizing a new planning horizon with the updated building models.

6.2 Optimization Testing

As previously noted, this work utilized the particle swarm algorithm implemented by Corbin et al. [24]. The implementation utilizes a multi-swarm approach where independent swarms are executed to reduce the likelihood of the algorithm settling in a local minimum. The PSO algorithm was chosen for its ability to accommodate "black box" building simulation function evaluators and was found to work well for optimization of building control setpoints. The algorithm was also utilized by Olivieri et al. to investigate similar building optimization problems with good result in [76]. May-Ostendorp et al. slightly modified the original implementation to allow for parallel execution of independent neighborhoods, reducing the overall time required for optimization [72]. In general, the work proposed in this research poses higher dimensional problems than those previously explored, and it was necessary to determine the optimization settings that were required to achieve satisfactory performance.

Numerous parameters can be considered in tuning the PSO algorithm, such as inertia weight, acceleration constants, and swarm sizes. No attempt was made in this work to exhaustively explore settings and achieve the most efficient convergence. Rather, it was assumed that higher dimensional problems are most likely to suffer from under exploration of the decision space, and attention was

given to determining appropriate swarm sizes for the problem dimensionality.

Preliminary testing was performed using homogenous portfolios of one to three buildings. Homogeneous portfolios were used since it seems logical that the opportunity for synergy should be small when buildings are identical. Furthermore, no synergy should be expected if the demand penalty is excluded from the optimization since building energy consumption is independent and there is no connection between operations. Cases including and excluding the demand penalty were explored.

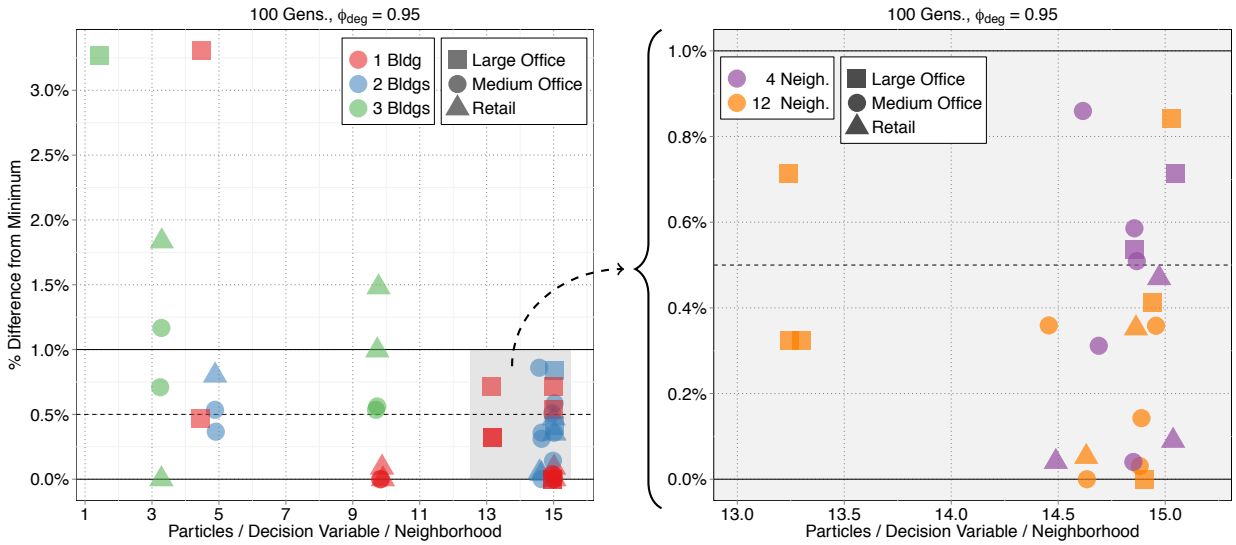

Figure 6.2: Preliminary optimization testing of swarm size.

The results of approximately 45 preliminary optimizations are shown in Figure 6.2. Since the number of decision variables differs between building models, it was desired to evaluate the swarm size with respect to problem dimensionality. Thus, the independent axis of Figure 6.2 is reported as the number of particles per decision variable for each independent neighborhood. The dependent axis is reported as the percent difference from the expected minimum. Due to the previously observed reliable performance on single building problems, the minimum objective function value observed for each single building case was used as the reference minimum cost. Simulation combinations were not performed exhaustively, and completed optimizations were used to inform the next round of optimizations.

Figure 6.2 shows that the single building large office and three-building optimizations came within 3.5% of the observed minimum value when one to five particles per decision variable (per neighborhood) were used. The two-building medium office and retail optimizations came within 1% of the minimum when using five particles. Medium office building and retail single building optimizations achieved their minimum cost around ten particles per decision variable per neighborhood, while the three-building portfolios were spread within 0.5% to 1.5% of their respective minimums. Three-building optimizations were abandoned for practical reasons after ten particles per decision variable, since runtimes were increasing significantly. The majority of testing was performed around 15 particles per decision variable per neighborhood since results appeared promising. All single building and two-building portfolio optimizations came within 1% of the expected minimum, with many less than 0.5% as well.

The region from 13 to 15 particles, between 0% and 1% is highlighted in the right panel of Figure 6.2. This excerpt shows that testing also included the evaluation of dispatching four independent neighborhoods versus 12. Acceptable results were observed for both numbers of neighborhoods, suggesting that four independent swarms with 15 particles per decision variable, per neighborhood, should be adequate for optimizing the two building portfolios of interest.

It is also worth briefly discussing the role of the inertia weight degradation factor (ϕ_{deg}) used in this PSO variant. The particle inertia weight (w) essentially scales the contribution of a particles' previous velocity to its current velocity, allowing the algorithm to be shifted between exploratory and locally refining natures [5]. In this implementation the inertia weight is described by an exponential function of the form $w(g) = \phi_{deg}^{g-1}$, where g is the optimization generation ($g \in [1,\ 2,\ 3, \ldots,\ g_{final}]$). For $0 < \phi_{deg} < 1$, the function is an exponential decay, as highlighted in Figure 6.3, with larger values of w representing stronger emphasis on global searching tendencies while smaller weights represent more focus on local searching. The previous optimizations were performed with $\phi_{deg} = 0.95$ which shows that the inertia weight is 0.5 after 13 generations, and near zero after 100. Considering this effect it seems logical that a greater number of particles would be required to adequately search the space, since many generations are primed to search locally.

If a larger degradation factor had been used, (e.g. $\phi_{deg} = 0.99$) the same results may have been achievable with less particles per decision variable since the algorithm would maintain its global searching properties longer. However, fewer particles does not necessarily translate into faster convergence since more generations may be required for the particles to move towards favorable regions. The performance of the PSO algorithm using 15 particles per decision variable and $\phi_{deg} =$

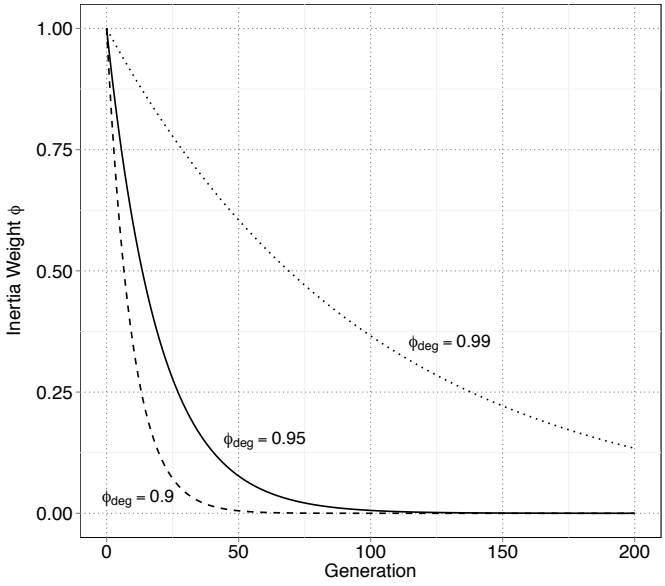

Figure 6.3: PSO inertia weight as a function of generation and degradation factor.

0.95 is also highlighted in Figure 6.4 by plotting the normalized cost of the best solution as a function of optimization generation. It is observed that the majority of savings are achieved in the first ten to 15 generations, presumably due to global search tendencies identifying near-optimal regions. Small refinements are then made as the algorithm transitions towards local searching. From these results it appears that executing PSO with $\phi_{deg} = 0.95$, four to 12 neighborhoods, and 15 particles per decision variable per neighborhood should achieve satisfactory results within 100 generations for the desired problems and dimensionality.

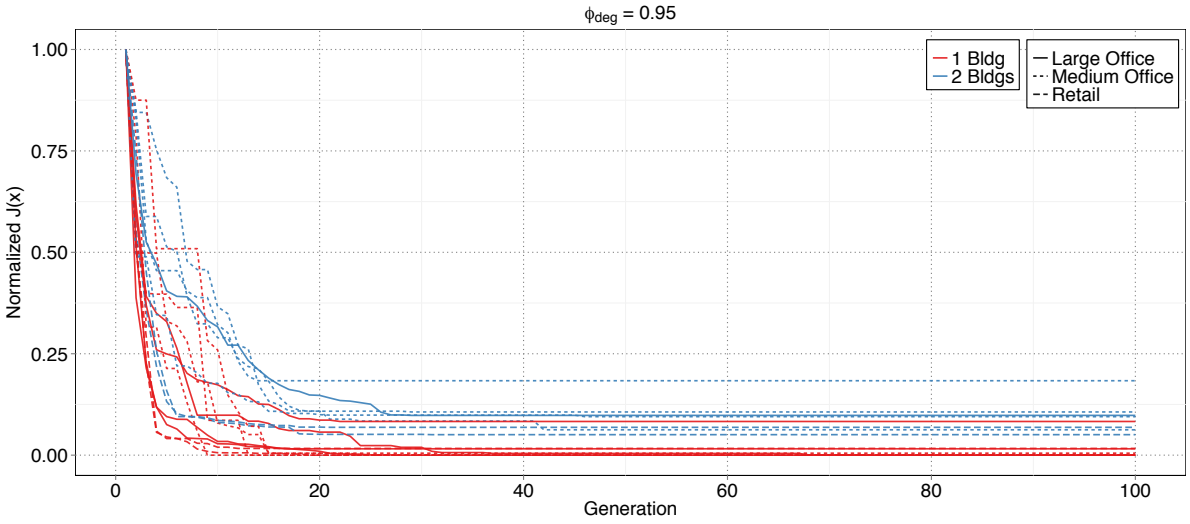

Figure 6.4: Normalized minimum cost as a function of optimization generation.

6.3 Preliminary Results

The first round of portfolio simulations was performed using all two-building combinations as shown in the simulation design of experiments in Table 6.1, where R denotes a retail building, M denotes a medium office, and L denotes a large office. All portfolios were simulated using the same high and low target demand limits as in the FR results presented in Chapter 5. In the case of heterogeneous portfolios, building multipliers were assigned so that the demand profiles of each portfolio member were of similar magnitude, in order to avoid the cost function being dominated by a single building. It was determined that 21 retail, eight medium office, and one large office building had demand profiles of similar magnitude.

The results for the high TDL case are summarized in Table 6.2, where the values are reported as the percent difference with respect to the aggregated individual optimization results. Positive values indicate the portfolio solution resulted in a higher final cost than the individual optimizations. Table 6.2 shows that the portfolio optimizations found solutions with nearly the same final cost as the aggregated individual results, although slightly higher. Further comparison of simulation results showed that strategies were practically identical between individual and portfolio optimizations, thus, no synergy was observed for the high TDL cases. After reviewing the target demand limit and

Table 6.1: Initial portfolio optimization design of simulation experiments.

	Low TDL			High TDL		
Bldg	R	M	L	R	M	L
R	RR	RM	RL	RR	RM	RL
M	-	MM	ML	-	MM	ML
L	-	-	LL	-	-	LL

simulation output, this result seems logical. The TDL was high enough that no demand penalties were incurred, which simplifies the optimization to only include energy expense and regulation revenue terms. In this scenario energy consumption and FR are independent between buildings, and the portfolio optimizations should find the same solutions as the single building optimizations. Overall, these results can be considered further verification that the optimization is performing adequately on the desired problems.

The results for the low TDL optimizations are summarized in Table 6.3. The homogeneous portfolios found virtually the same final cost as the aggregated individual optimizations. Further investigation showed that the target demand limit was set below the minimum achievable portfolio demand, and even the lowest obtainable demand incurred a penalty. This low target demand, in essence, drove all buildings in the portfolio toward maximum demand reducing strategies. When portfolio members are identical, the solutions that achieve minimum peak demand are also identical and no diversity exists among the portfolio (i.e. buildings peak during the same hour).

However, results for the heterogeneous portfolios look more promising. Table 6.3 shows that OPT cases for the RL and ML portfolios achieved 2.45% and 2.95% lower final cost, respectively, than the aggregated individual results. Similarly, the OPT+FR RL and ML cases achieved 1.87% and 2.2% savings over the aggregated individual solutions. Although savings were observed, further examination of simulation results showed that setpoint strategies were virtually the same between portfolio and aggregated individual results. The low demand target, once again, forced buildings towards demand minimizing strategies. The difference in the heterogenous RL and ML cases is that

Table 6.2: Portfolio % difference from aggregated individual optimizations - high TDL.

Type	Case	NSU	OPT	NSU+FR	OPT+FR	Unit
Homogeneous Portfolio	RR	0.00	0.00	0.00	0.00	%
	MM	0.00	0.14	0.00	0.66	%
	LL	0.00	0.84	0.00	1.25*	%
Heterogeneous Portfolio	RM	0.00	0.00	0.00	0.00	%
	RL	0.00	0.24	0.00	0.43	%
	ML	0.00	0.33	0.00	0.91	%

*Only completed 20 generations.

Table 6.3: Portfolio % difference from aggregated individual optimizations - low TDL.

Type	Case	NSU	OPT	NSU+FR	OPT+FR	Unit
Homogeneous Portfolio	RR	0.00	0.26	0.00	0.26	%
	MM	0.00	0.33	0.00	0.91	%
	LL	0.00	0.41	0.00	2.39*	%
Heterogeneous Portfolio	RM	0.00	0.22	0.00	0.68	%
	RL	0.00	-2.45	-0.18	-1.87	%
	ML	0.00	-2.95	-0.18	-2.20	%

*Only completed 20 generations.

the retail building and medium office building peak at hour ending 4:00 PM, while the large office building peaks at hour ending 1:00 PM. This diversity results in lower overall peak for the building portfolios and reduces the demand penalty by $160 to $200, depending on the case. Although the peak hours occur coincidently among the NSU cases, the RL and ML NSU+FR cases were able to achieve slightly more regulation revenue due to diversity among other on-peak hours. This resulted in 0.18% savings over the aggregated individual optimizations. It is also noted that the magnitude of savings between OPT and OPT+FR cases cannot be directly compared since they are computed with respect to different individual optimization cases.

Although the low TDL results for heterogeneous portfolios did show savings over the individually optimized cases, the buildings were essentially coerced by the demand limit to the same solutions as the individual optimizations. The additional savings due to demand diversity may be significant, however, it is a stretch to label them as synergistic effects achieved through portfolio optimization. Simply recomputing the demand penalty to accommodate diversity among the individual optimizations would have determined similar savings. Noteworthy synergistic effect would be observed if the portfolio optimization chose to operate buildings in a state that is suboptimal from an individualistic perspective, but that leads to greater overall benefit. In order for such a scenario to be observed it appears that buildings must have more operating freedom and not be forced into peak minimizing strategies. Prior to optimization it is difficult to know the magnitude of achievable demand reductions, thus the optimization results of this chapter were used in designing additional simulation studies. Chapters 7, 8, and 9 present further portfolio simulation case studies with higher target demand limits that allow the desired demand reductions to potentially be achieved through a variety of control scenarios.

Chapter 7

Portfolio Case 1: Retail and Large Office Buildings

The results of the initial portfolio simulation study determined that the large office building was capable of reducing peak demand by a maximum of 808 kW and the retail building was able to achieve a maximum demand reduction of about 6.6 kW. It seems possible that synergistic effects may exist when buildings are not steered toward maximum demand reducing strategies and have flexibility to trade off responsibilities. To test this hypothesis a portfolio was constructed consisting of 120 retail buildings and a single large office building. This ratio was chosen so that the building types under control in the optimization have equal ability to reduce demand (i.e. 120 retail buildings can shed approximately the same amount of load as one large office building). The portfolio-level target demand limit was set to encourage demand reductions without forcing all buildings into maximum load reducing strategies. Figure 7.1 illustrates the targeted portfolio demand in relation to the NSU peak and absolute minimum achievable demand. As a whole, the portfolio is seeking to shed 808 kW of peak demand. Due to the portfolio construction it is possible to achieve this load reduction using only one building type or by splitting the reductions between different types.

The portfolio optimization determines the appropriate operation of each portfolio member to meet portfolio wide objectives. In order to illuminate synergistic effect, the portfolio optimizations were compared to the aggregation of individually optimized portfolio members. Optimizing portfolio members individually requires dividing the portfolio TDL into separate smaller limits for each building type. There are an infinite number of ways to achieve this task, however three scenarios have been selected for evaluation in this work. The three scenarios are shown in Figure 7.1.

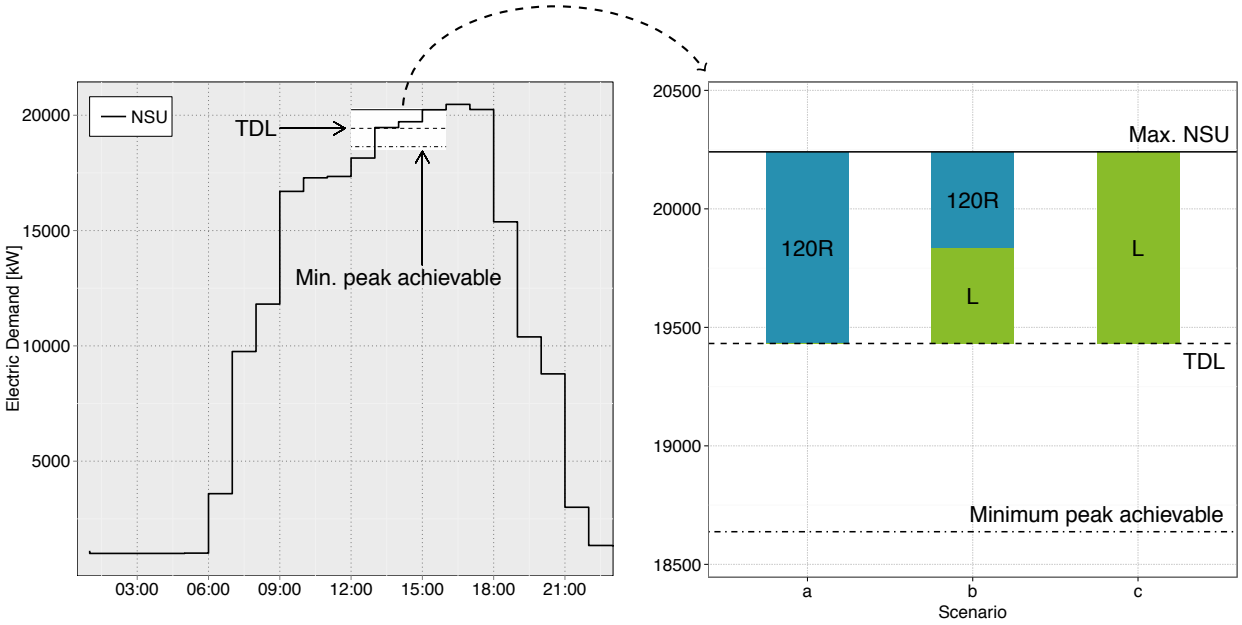

Figure 7.1: Case 1 scenarios and portfolio demand limits.

Scenario "a" represents the case where the retail building receives a low target demand limit such that it would perform the entire load reduction desired by the portfolio (i.e. 808 kW). In scenario "a" the large office building consequently receives a high target demand limit such that reducing peak demand is of no concern. Scenario "b" represents the case where the retail building and large office building share the desired portfolio load reduction equally. Scenario "c" represents the case where the large office would be responsible for the entire portfolio demand reduction, while the retail building receives a high target demand and need not reduce load. All three scenarios lead to the same portfolio optimization, however, the outcome of the differing individual optimizations may have significant implications on whether synergistic effect is observed. In the following results, "OPT" and "NSU" are used to distinguish between the optimal and night set-up (or set-back) operations, respectively. A lower case "s" and "p" are also used to denote single optimization results and portfolio optimization results, respectively.

Table 7.1: Case 1 scenario target demand limits.

	Target Demand Limit			
Scenario	Rs	Ls	(RL)p	Units
a	16 419	3013	19 432	kW
b	16 826	2606	19 432	kW
c	17 227	2205	19 432	kW

7.1 Case 1a: Retail Building Requires Large Demand Reduction

As earlier noted, scenario "a" represents the situation where the desired portfolio demand reduction is entirely attributable to the desired reduction of the retail building. This scenario could occur when the buildings set their monthly peak demand on different days during the month. It seems possible that this could occur if the buildings are conditioned differently throughout the month, such as during weekends or holidays. To simulate this scenario the retail building target demand limit was set at 16 419 kW, requiring a load reduction of 808 kW to avoid incurring a demand penalty. The target demand limit for the individual large office optimization was set at 3013 kW, which is nearly equivalent to the NSU peak, resulting in no demand penalty for NSU operation.

7.1.1 No Frequency Regulation

The retail building optimization results for scenario "a" are highlighted in Figure 7.2, with the "on-peak" period extending from hour ending 1:00 PM to hour ending 4:00 PM. To achieve the desired demand reduction, the OPT Rs case chose an early startup with significant precooling. This resulted in the consumption of 56.3 MWh (28%) more energy and $1315 (12%) more energy expense than the NSU Rs case, with the reward of avoiding expensive demand penalties. However, when the portfolio perspective was considered in the OPT Rp case, the optimizer chose the NSU strategy, performing no load shifting through the retail buildings.

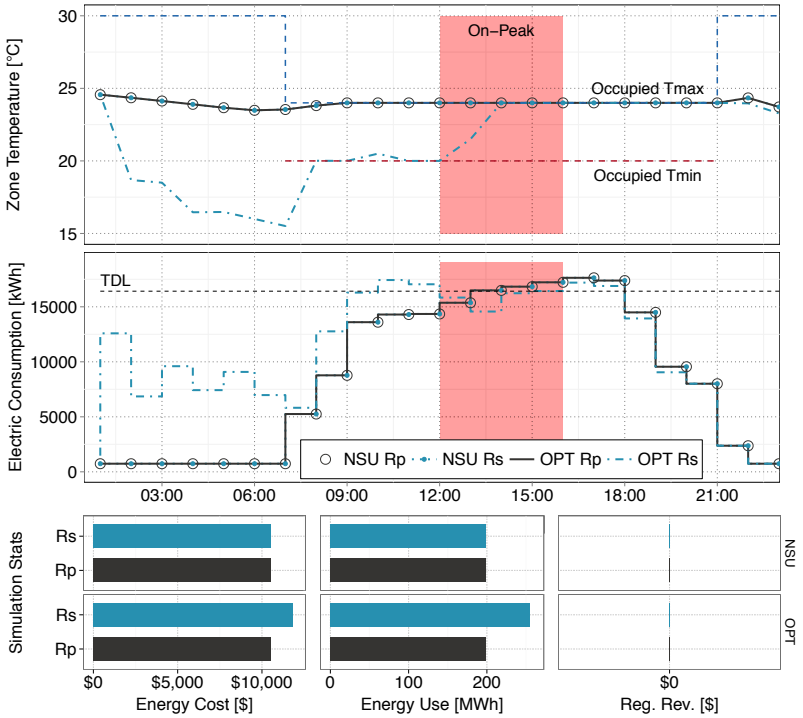

Figure 7.2: Case 1a retail (x120) optimization results.

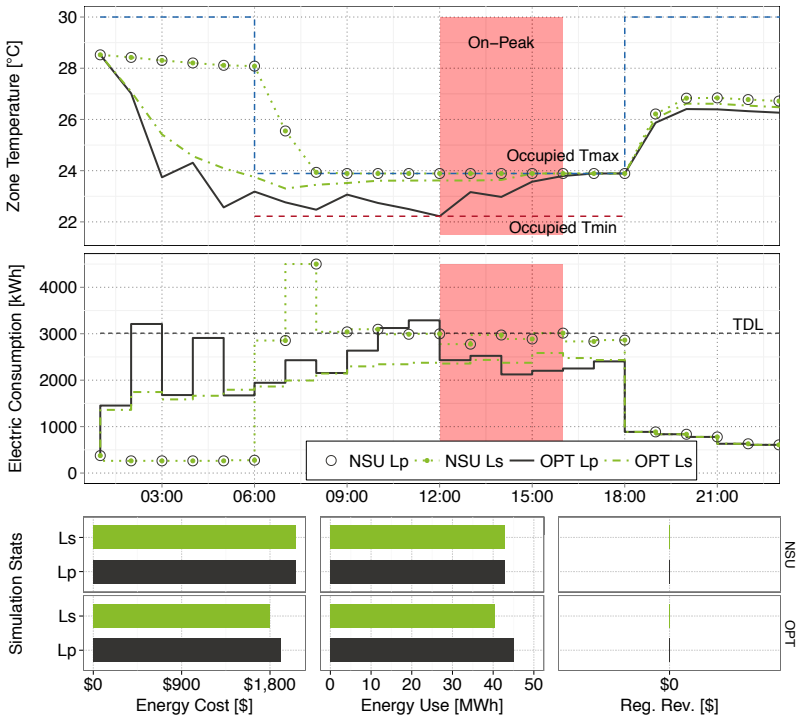

Figure 7.3: Case 1a large office (x1) optimization results.

The large office building results for scenario "a" are provided in Figure 7.3. When individually optimized (OPT Ls), light precooling was performed before high priced afternoon hours which generated savings from economizing, efficiency increases, and energy arbitrage (i.e. shifting energy use to lower priced hours). The precooling was incentivized purely by energy price, and resulted in a "voluntary" peak demand reduction of 428 kW. When optimized as a portfolio (OPT Lp), significantly more precooling was observed. The OPT Lp case reduced on-peak demand by 811 kW, achieving the reduction desired by the entire portfolio. Comparing the energy cost between the OPT Ls and OPT Lp solutions shows that greater energy expense is incurred as a result of the increased precooling, which would be a suboptimal solution for the large office building outside of the portfolio.

Further analysis of the portfolio-level results provides insight into the benefit created by using the large office building to perform all of the load reduction. Figure 7.4 highlights the portfolio energy consumption, optimizer cost $J(\vec{x})$, and the percent savings for each optimization case. Rs+Ls represents the combination of individually optimized retail and large office buildings, while (RL)p denotes results from the portfolio optimization. Percent savings are reported with respect to the Rs+Ls NSU case (without frequency regulation).

First, it is noted that the NSU Rs+Ls case is equivalent to the NSU (RL)p case. This implies that no natural diversity exists through simple aggregation of the retail and large office buildings under NSU operation, since both peak during hour ending 4:00 PM. Second, the top panel shows that providing the 808 kW demand reduction using the retail building (OPT Rs+Ls) expends 51.7 MWh more energy than when all shifting is performed by the large office building ((RL)p). This extra energy consumption translates into $1200 in additional energy expense. In the portfolio optimization the optimizer can see that the large office building can voluntarily shed over half the required load, and determined that the large office could further reduce demand at lower cost than the retail buildings. Overall the portfolio optimization increased the percent savings by 7.5 percentage points compared to the aggregated individual optimizations. Cost function values

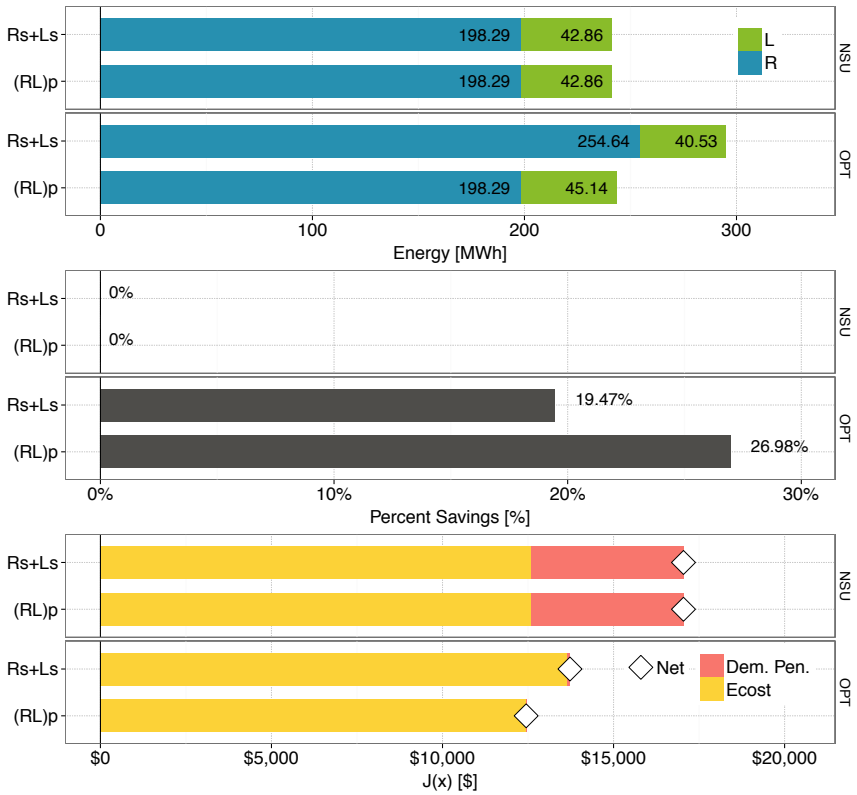

Figure 7.4: Case 1a aggregated individual and portfolio optimization results.

Table 7.2: Case 1a aggregated individual and portfolio optimization results.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	12 597	13 647	12 597	12 445
Energy Use [kWh]	241 146	295 174	241 146	243 430
Demand Penalty [$]	4447.19	78.61	4447.16	0
$J(\vec{x})$ [$]	17 044	13 726	17 044	12 445
% Difference	0.00%	-19.47%	0.00%	-26.98%

are highlighted in Table 7.2 for further reference.

7.1.2 With Frequency Regulation

Scenario "a" was repeated including frequency regulation to evaluate the impact of potential ancillary service revenues on the optimization. The large office was assumed to be available to provide FR from 9:00 AM to 5:00 PM. The retail building was defined as having no ability to perform frequency regulation through changes in HVAC operation since the RTU fan and DX coil are single speed and unable to modulate smoothly while following regulation dispatch. Thus, by definition, the OPT Rs results with frequency regulation are equivalent to those shown in the previous section. The results for the retail building are shown in Figure 7.5. Early start-up with significant precooling was observed in the OPT Rs case in order to achieve the target demand of 808 kW below NSU peak. The OPT Rp case resulted in NSU operation for the retail building as in the previous section due to the significant energy consumption associated with load shifting in the retail building.

The large office building results for scenario "a" including frequency regulation are shown in Figure 7.6. The NSU Ls case was able to perform $206 in regulation services through modulation of the chilled water setpoint, while the portfolio NSU case (NSU Lp) was able to perform $273 in regulation revenue. During on-peak hours, the regulation capability was often constrained by the target demand limit, since the requirement for symmetric regulation was enforced. In the presence of steep demand penalties, the FR estimation essentially attempts to provide FR up to the TDL. Since the retail building does not provide FR or operate at its peak demand for the entire on-peak period, additional space was available while adhering to the same demand limit. When the portfolio perspective was considered, the large office building was able to utilize the additional space beneath the TDL during hours when the retail building was below its peak and generate the additional $67 (33%) in regulation revenue for the NSU Lp case.

The OPT Ls case started cooling the building five hours before the NSU Ls case and kept temperatures near the middle of the occupied temperature temperature range, to reduce energy

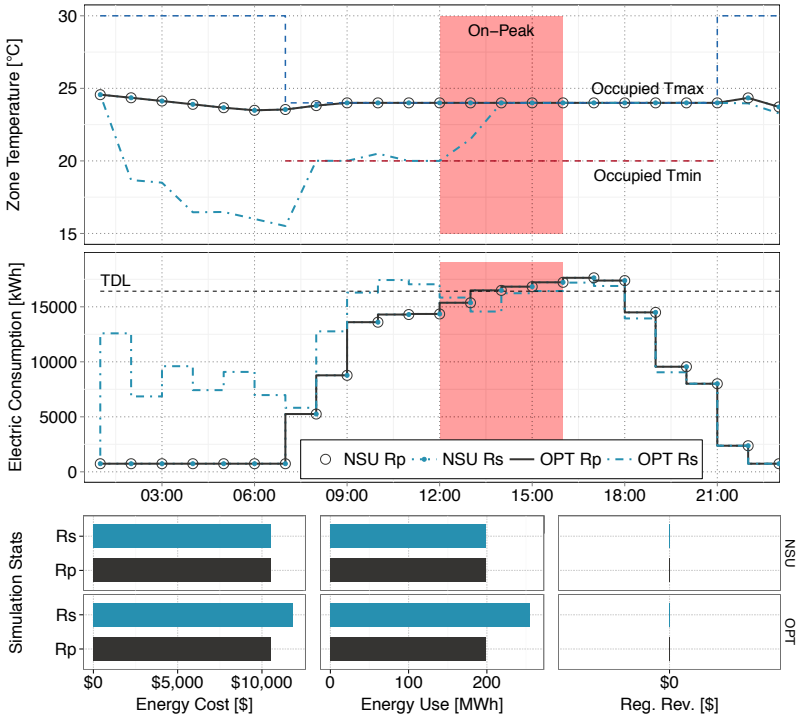

Figure 7.5: Case 1a retail (x120) optimization results including FR.

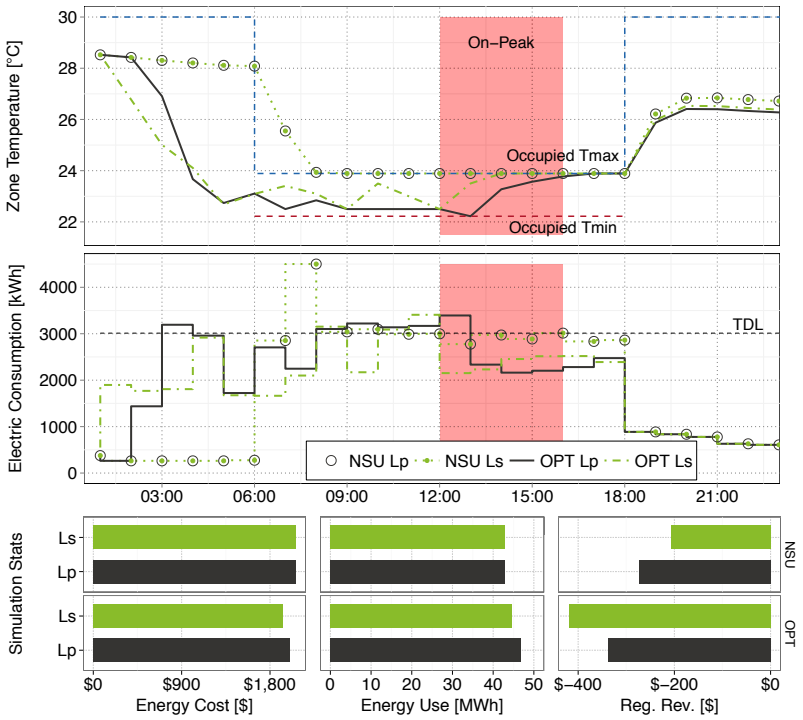

Figure 7.6: Case 1a large office (x1) optimization results including FR.

use during high-priced afternoon hours. The OPT Ls case was able to generate $418 in total regulation revenue and reduce energy cost by $134 at the expense of consuming an additional 1761 kWh more than the NSU Ls case. In the OPT Lp case the optimizer chose to keep the zone temperatures slightly lower during the hours preceding the peak period in order to achieve greater demand reductions and meet the portfolio TDL. Although the entire 808 kW demand reduction is performed by the large office building it is still able to achieve $337 of regulation revenue.

Figure 7.7 highlights the portfolio-level results in terms of energy consumption, percent savings, and optimizer cost. The top panel shows that the NSU energy consumption was identical between the single building and portfolio optimizations, and that the OPT Rs+Ls case consumed 54.2 MWh more than the OPT (RL)p case to meet the target demand limit. Although the NSU cases were able to generate regulation revenue, the savings are relatively small at 1.2% to 1.6%. This is mostly due to the large proportion of retail buildings in the portfolio that contribute significantly to the energy cost and demand penalty terms in the cost function, but are unable to participate in regulation.

Similarly the contribution of FR to the overall savings in the OPT cases was relatively small as well. The bottom panel shows that the OPT (RL)p case was able to eliminate the demand penalty at lower energy cost than the OPT Rs+Ls case. Overall the portfolio optimization achieved 7.3 percentage points more savings than the individually optimized buildings. Detailed cost function values are provided in Table 7.3.

7.2 Case 1b: Buildings Require Similar Demand Reductions

Scenario "b" represents the situation where individually the large office and retail building contribute equally to the desired portfolio demand reduction (i.e. 404 kW each). It is conceivable to imagine that this scenario may arise when both buildings are attempting to operate under previously set peak demands, and must shed similar amounts of load to avoid setting a new peak. This scenario was simulated by setting the Rs target demand at 16 826 kW and the Ls target demand at 2606 kW. The portfolio TDL remains at 19 432 kW thus portfolio optimization solutions

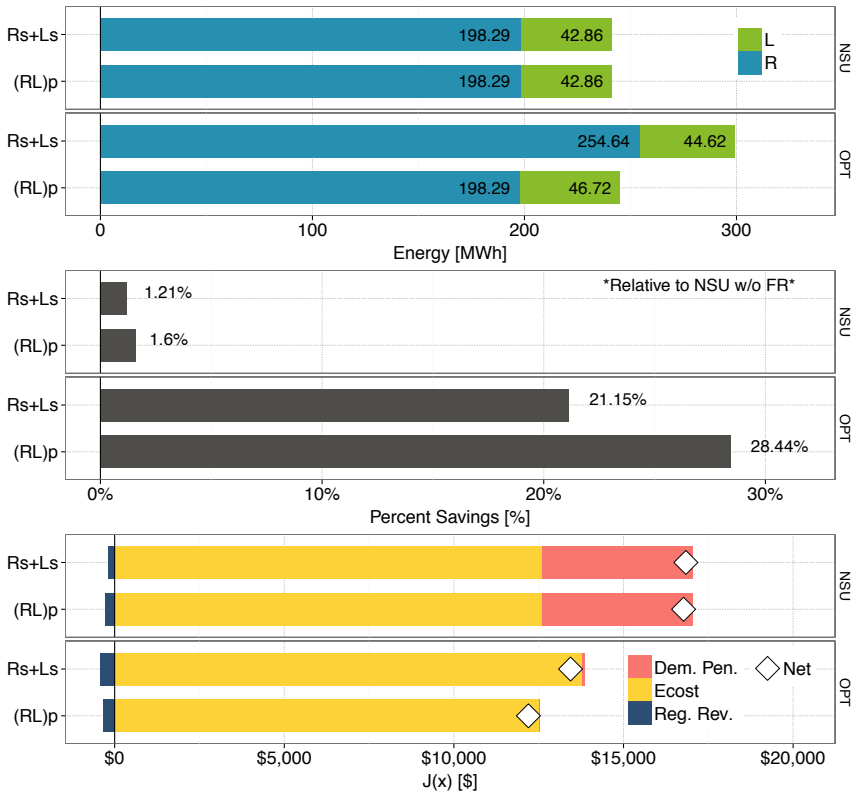

Figure 7.7: Case 1a aggregated individual and portfolio optimization results with FR.

Table 7.3: Case 1a aggregated individual and portfolio optimization results with FR.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	418.12	273.30	337.44
Energy Cost [$]	12 597	13 778	12 597	12 535
Energy Use [kWh]	241 146	299 260	241 146	245 002
Demand Penalty [$]	4447.19	78.61	4447.16	0
$J(\vec{x})$ [$]	16 838	13 439	16 771	12 197
% Difference	-1.21%	-21.15%	-1.60%	-28.44%

are identical to scenario "a".

7.2.1 No Frequency Regulation

The retail optimization results for scenario "b," excluding frequency regulation, are shown in Figure 7.8. The OPT Rs case started conditioning two hours before typical NSU operation and cools to the lower setpoint boundary of 20 °C immediately preceding the on-peak period to achieve the 404 kW peak load reduction. The precooling is less extreme than in scenario "a," as expected for the higher TDL. The load shifting consumes 15.4 MWh more than the NSU case, resulting in $485 of additional energy expense. As previously observed, the OPT Rp solution follows the NSU setpoint strategy.

Results for the large office optimizations are shown in Figure 7.9. Early start-up and light precooling were observed in the OPT Ls case to achieve the desired 404 kW demand reduction and avoid energy consumption during high priced hours. The OPT Ls case is able to save 1403 kWh of energy through economizing and efficiency improvements, and achieve an overall energy cost savings of $258. The OPT Lp case precools near the lower temperature bound of 22.2 °C preceding the on-peak period in order to achieve greater demand reductions and meet the 808 kW portfolio-level demand reduction. Reducing additional load in accommodation of the retail buildings requires consuming 3687 kWh more than the OPT Ls case, resulting in $107 in additional energy expense.

Figure 7.10 summarizes the aggregate results for scenario "b." The NSU cases are identical in energy consumption, percent savings, and optimizer cost implying that no natural diversity exists under NSU operation. By performing all load shifting via the large office building, the (RL)p OPT case is able to reduce energy consumption by 11.7 MWh (4.6%) and energy expense by $378 (2.9%) over the individually optimized case (OPT Rs+Ls), resulting in a 2.2 percentage point increase in total percent savings. Cost function numbers for the aggregated individual and portfolio optimizations are provided in Table 7.4.

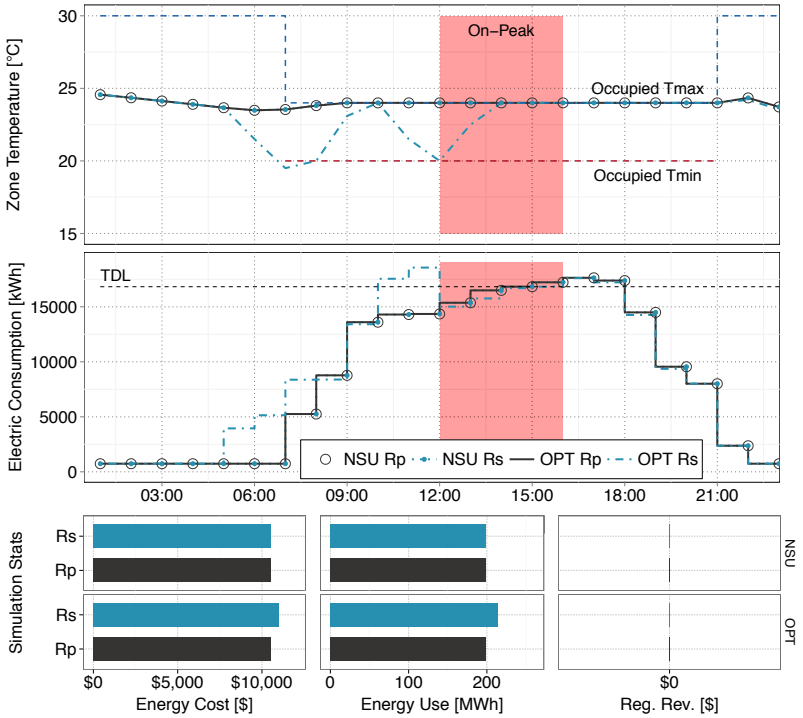

Figure 7.8: Case 1b retail (x120) optimization results.

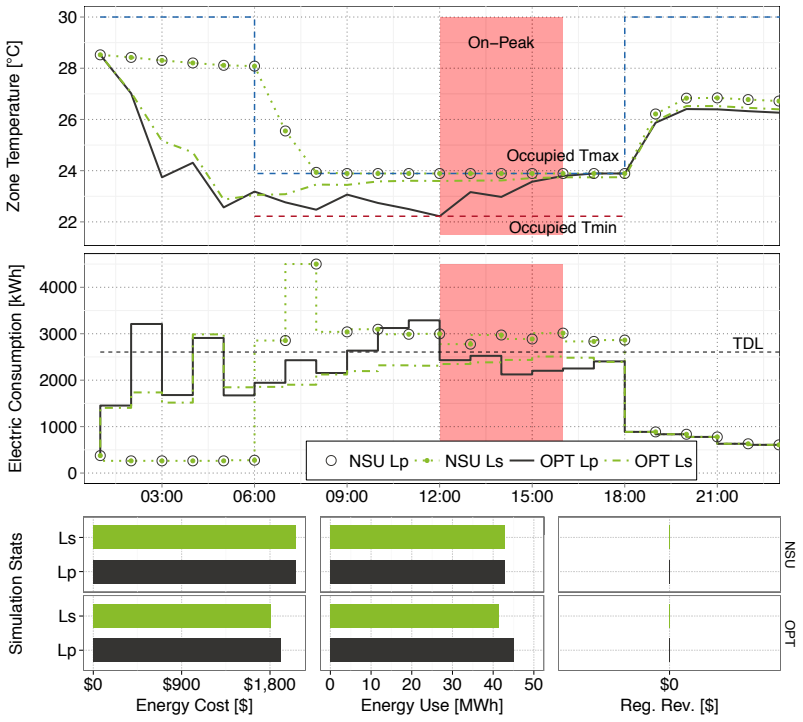

Figure 7.9: Case 1b large office (x1) optimization results.

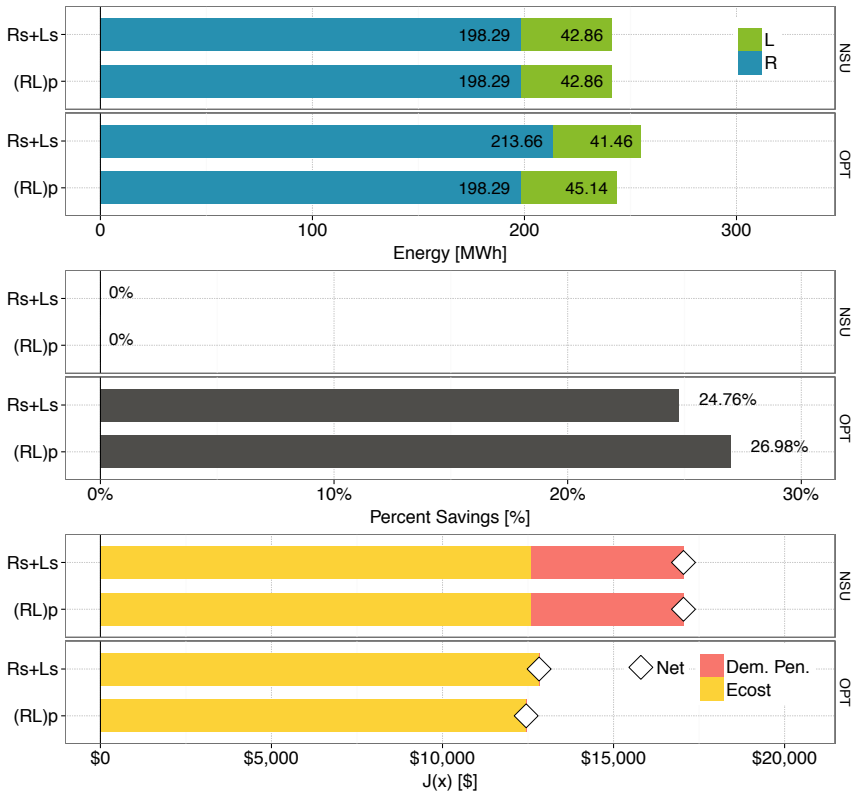

Figure 7.10: Case 1b aggregated individual and portfolio optimization results.

Table 7.4: Case 1b aggregated individual and portfolio optimization results.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	12 597	12 823	12 597	12 445
Energy Use [kWh]	241 146	255 121	241 146	243 430
Demand Penalty [$]	4447.19	0	4447.16	0
$J(\vec{x})$ [$]	17 044	12 823	17 044	12 445
% Difference	0.00%	-24.76%	0.00%	-26.98%

7.2.2 With Frequency Regulation

The scenario "b" retail building results including frequency regulation are shown in Figure 7.11. Since the retail building is unable to provide FR by definition, the results are identical to the non-FR case previously shown in Figure 7.8.

The large office results for scenario "b" are highlighted in Figure 7.12. As previously observed, the NSU Ls case is able to generate $206 of regulation revenue, while the NSU Lp case is able to generate $273 due to the additional space and diversity created through aggregation. When optimized individually (OPT Ls), the large office building precools and maintains a zone temperature near 22.5 °C preceding the on-peak period. The stored thermal energy is then discharged during the peak period to achieve the necessary 404 kW demand reduction. A closer look at the demand profiles reveals the optimizer chose to shed 651 kW of peak demand in order to create more space for providing frequency regulation. The OPT Ls case was able to generate $384 of regulation regulation revenue while meeting the 2606 kW TDL, nearly doubling the revenue generated in the NSU Ls case. The increased precooling requires consuming 2.7 MWh more energy, costing an additional $119. In the portfolio optimization (OPT Lp), the large office takes on the responsibility of the entire 808 kW demand reduction which results in increased energy consumption and reduced regulation revenue when compared to the OPT Ls case.

The portfolio results for scenario "b" including regulation are summarized in Figure 7.13. Overall the results are similar to the non-FR cases since the regulation revenue generated by the office building tends to be small relative to the demand penalties and energy cost of the entire portfolio. Adding regulation to the NSU cases results in 1.2% to 1.6% savings over the non-FR NSU cases. The portfolio optimization increased the total percent savings by 2.2 percentage points over the individually optimized cases due to energy cost and energy savings associated with the more efficient load shifting performed by the large office building. The cost function breakdown is numerically summarized in Table 7.5 for reference.

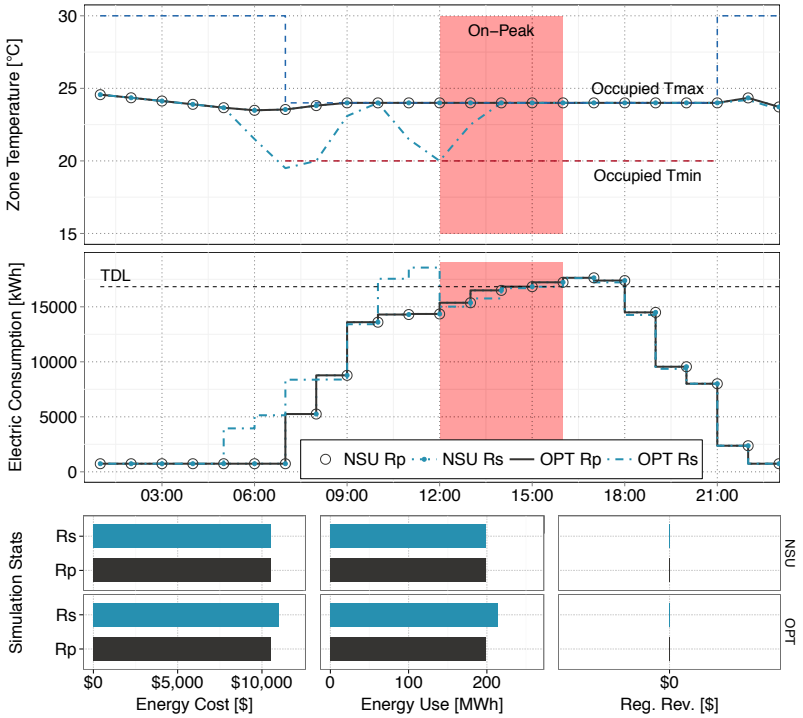

Figure 7.11: Case 1b retail (x120) optimization results including FR.

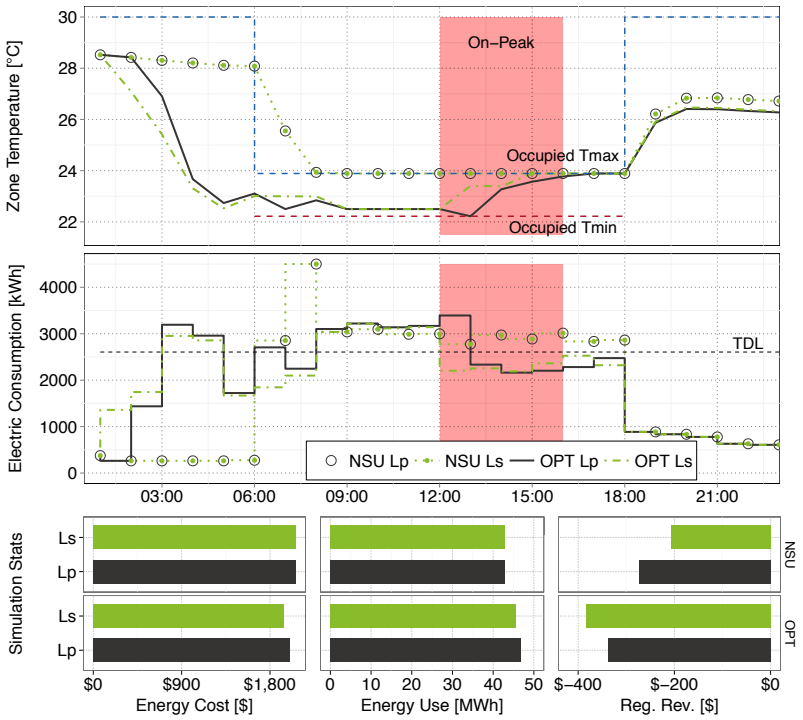

Figure 7.12: Case 1b large office (x1) optimization results including FR.

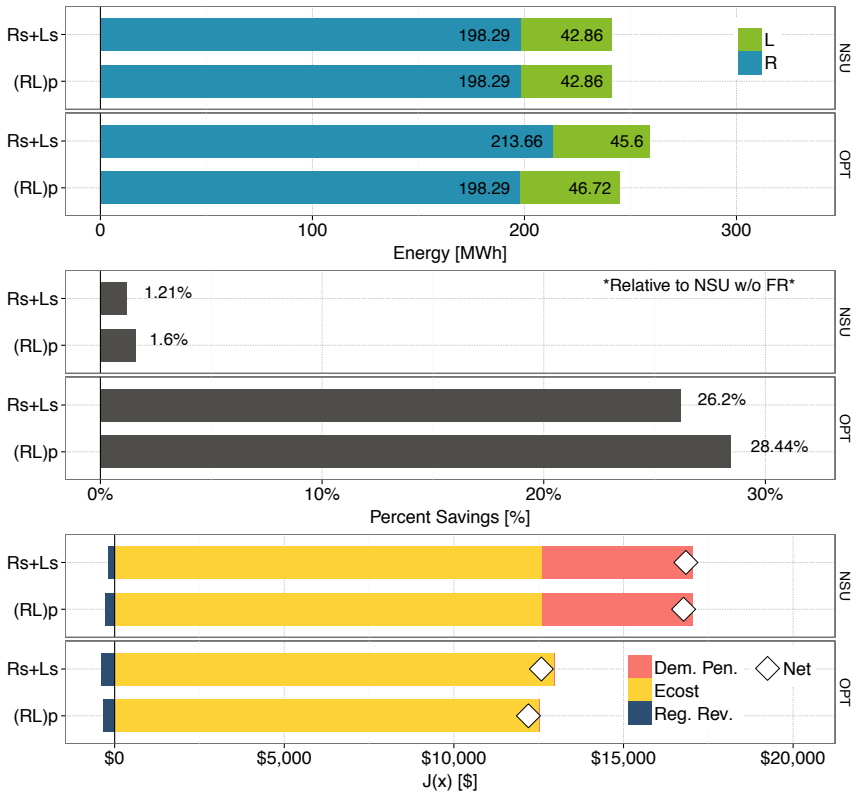

Figure 7.13: Case 1b aggregated individual and portfolio optimization results with FR.

Table 7.5: Case 1b aggregated individual and portfolio optimization results with FR.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	383.66	273.30	337.44
Energy Cost [$]	12 597	12 963	12 597	12 535
Energy Use [kWh]	241 146	259 266	241 146	245 002
Demand Penalty [$]	4447.19	0	4447.16	0
$J(\vec{x})$ [$]	16 838	12 579	16 771	12 197
% Difference	-1.21%	-26.20%	-1.60%	-28.44%

7.3 Case 1c: Large Office Requires Large Demand Reduction

Scenario "c" represents the situation where individually the large office desires to reduce demand by 808 kW, while the retail building has no need to reduce load. This scenario might also arise if the building set their monthly peak on separate days as in scenario "a." This scenario was simulated by setting the large office TDL at 2205 kW and the retail TDL at 17 227 kW. The portfolio TDL remains at 19 432 kW, thus, portfolio optimization results are identical to scenarios "a" and "b".

7.3.1 No Frequency Regulation

The optimization results for the retail building in scenario "c" are shown in Figure 7.14. Due to the high TDL, the NSU Rs case incurs no demand penalty, and is only incentivized by real-time energy prices to shift consumption. The optimizer determined that no beneficial opportunities for energy arbitrage existed and that NSU operation was optimal. As previously observed, the portfolio results determined load shifting using the retail buildings was not effective and NSU operation was chosen as well.

The large office optimization results for scenario "c" are highlighted in Figure 7.15. The OPT Ls case performs significant precooling and maintains the lower zone temperature boundary of 22.2 °C preceding the peak period to achieve the desired load reduction. The strategy chosen in the OPT Lp case is similar, however, a slightly warmer temperature was able to be maintained preceding the peak period to achieve the same demand reductions. Further analysis of the electric profile shows that the large office building was able to keep a higher demand during hours ending 1:00 PM and 2:00 PM in the portfolio optimization due to diversity created through aggregation with the retail building. This resulted in the Lp OPT case being able to achieve the 808 kW demand reduction using 4439 kWh (1.8%) less than the OPT Rs+Ls case, saving approximately $75 (0.6%) in energy cost for the day.

The portfolio results for scenario "c" are summarized in Figure 7.16. The NSU results were

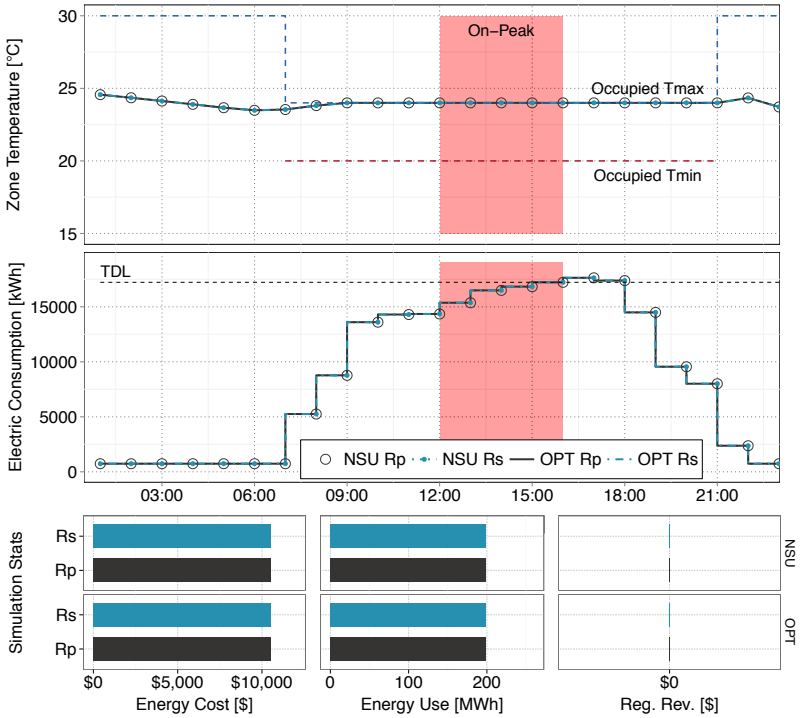

Figure 7.14: Case 1c retail (x120) optimization results.

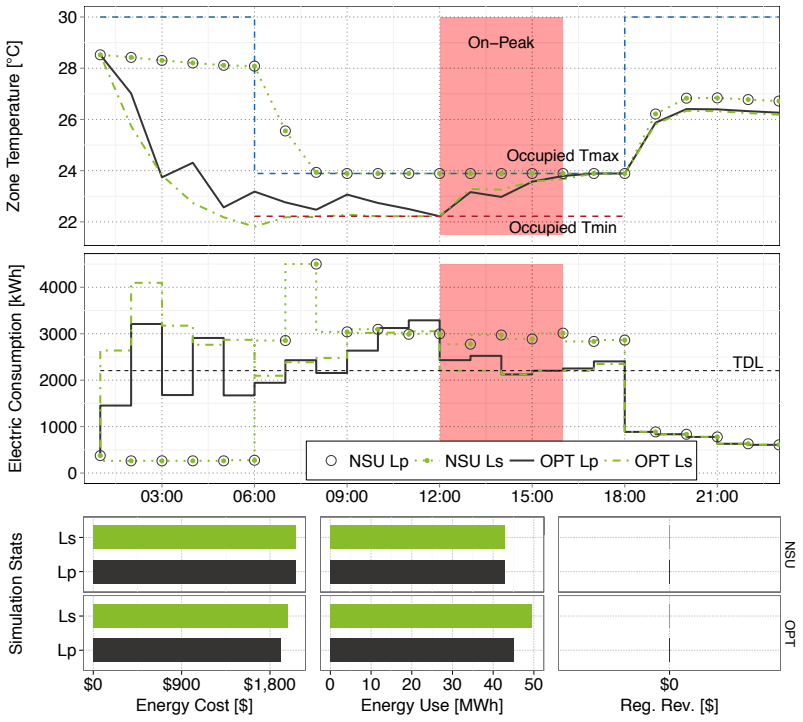

Figure 7.15: Case 1c large office (x1) optimization results.

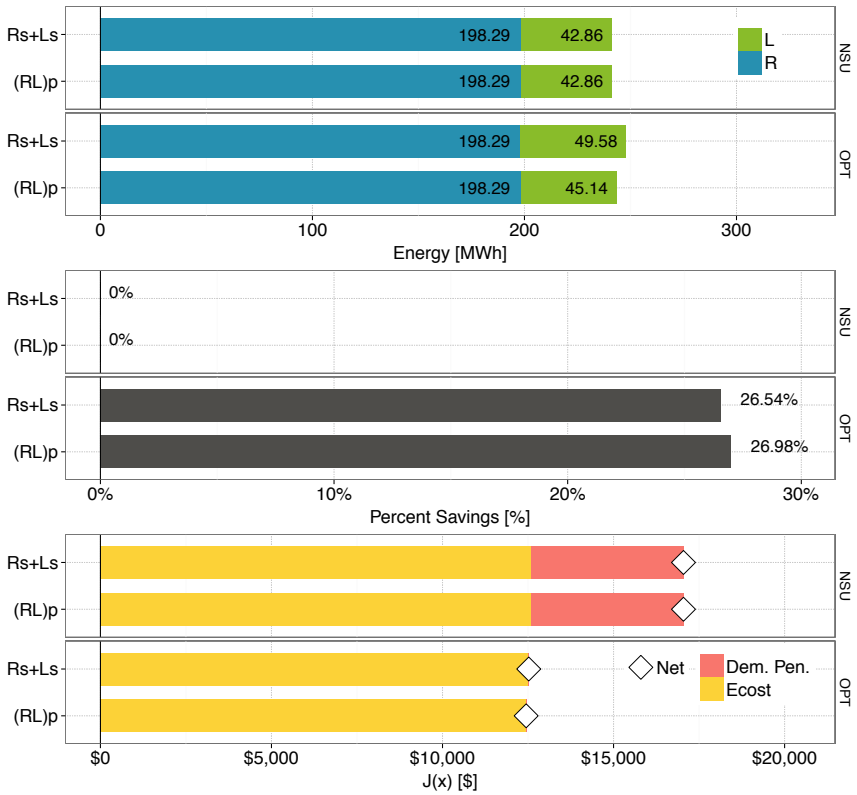

Figure 7.16: Case 1c aggregated individual and portfolio optimization results.

Table 7.6: Case 1c aggregated individual and portfolio optimization results.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	12 597	12 520	12 597	12 445
Energy Use [kWh]	241 146	247 869	241 146	243 430
Demand Penalty [$]	4447.19	0.77	4447.16	0
$J(\vec{x})$ [$]	17 044	12 521	17 044	12 445
% Difference	0.00%	-26.54%	0%	-26.98%

identical in energy consumption and optimizer cost implying synergy does not naturally exist through aggregation under NSU operation for these specific buildings. Slight energy savings was observed in the portfolio case as previously noted, however the result is only a 0.4 percentage point increase in total percent savings.

7.3.2 With Frequency Regulation

Scenario "c" results for the retail building including regulation are shown in Figure 7.17. As in the previous results, the retail building does not perform regulation and NSU operation was determined to be optimal in both individual and portfolio optimizations. The large office building results for scenario "c" are shown in Figure 7.18, and exhibit precooling strategies similar to those observed in the non-FR results. The NSU Lp case was once again able to generate $67 more regulation revenue than the NSU Ls case due to the retail building not performing regulation and operating below the TDL during several on-peak hours. Considering the OPT cases, the portfolio optimization (OPT Lp) was able to achieve the desired demand reduction while consuming 3 MWh less energy and providing $86 more in regulation revenue.

The portfolio results for scenario "c" are summarized in Figure 7.19 and show that a minor savings increase of 0.6 percentage points was observed when buildings were optimized as a portfolio. As with the previous Case 1 examples, regulation revenue was a fairly small contributor to the overall cost function.

7.4 Case 1 Summary

Comparing the results across scenarios "a", "b", and "c" for Case 1 highlights the dependence of synergistic effect on the conditions enforced during individual optimizations. Table 7.8 summarizes the achieved percent savings by the OPT cases for the investigated scenarios. The NSU Rs+Ls case was identical for all scenarios, which allows for fair comparison between percent savings across scenarios. The portfolio results determined that for this specific portfolio construc-

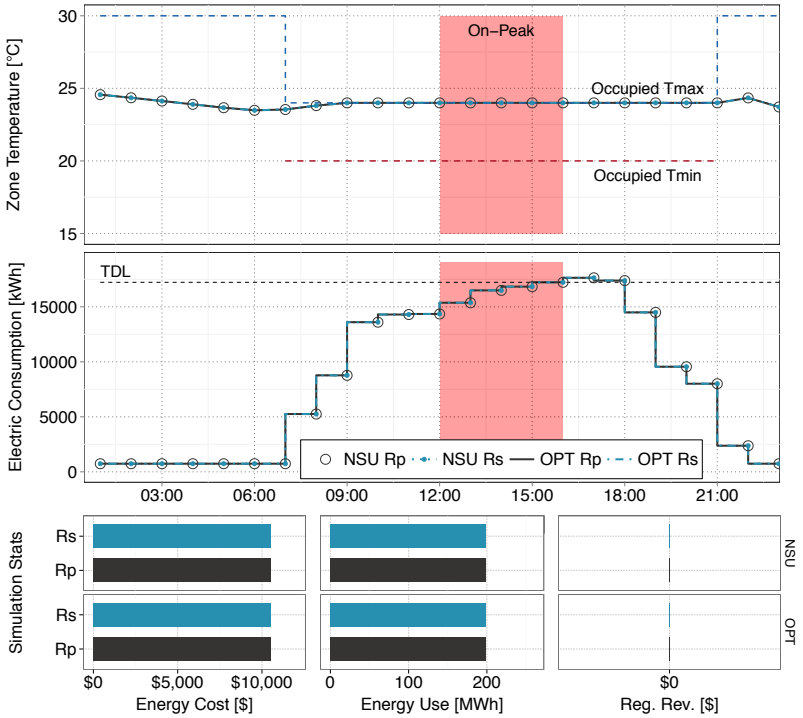

Figure 7.17: Case 1c retail (x120) optimization results including FR.

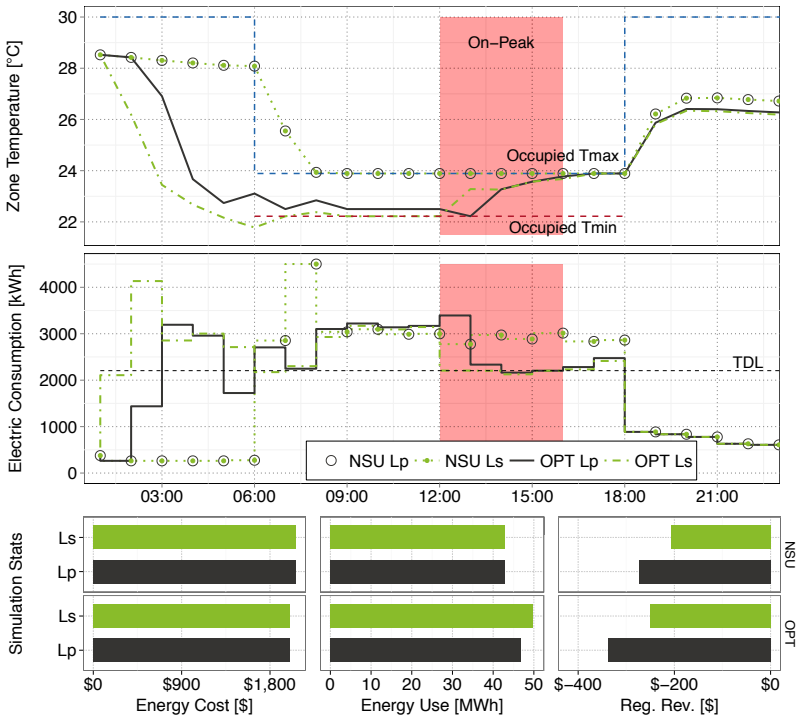

Figure 7.18: Case 1c large office (x1) optimization results including FR.

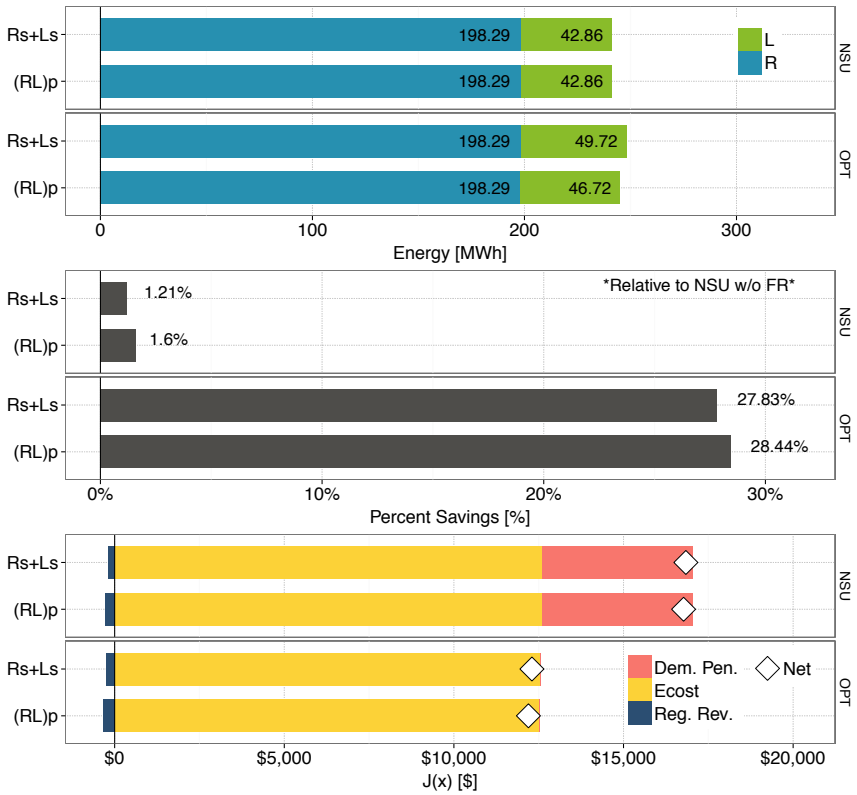

Figure 7.19: Case 1c aggregated individual and portfolio optimization results with FR.

Table 7.7: Case 1c aggregated individual and portfolio optimization results with FR.

	Rs+Ls		(RL)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	251.16	273.30	337.44
Energy Cost [$]	12 597	12 542	12 597	12 535
Energy Use [kWh]	241 146	248 001	241 146	245 002
Demand Penalty [$]	4447.19	9.35	4447.16	0
$J(\vec{x})$ [$]	16 838	12 300	16 771	12 197
% Difference	-1.21%	-27.83%	-1.60%	-28.44%

tion (i.e. 120 retail, and one large office) the desired load reductions were optimally achieved when the load shift was performed entirely by the large office building. The individual demand limits set in scenario "c" lead to virtually the same solution as the portfolio optimization, thus, only 0.4 to 0.6 percentage points of additional savings were observed when the buildings were simultaneously optimized. However, scenario "a" was quite the opposite of the optimal portfolio solution, and therefore the simultaneous optimization achieved seven additional percentage points in total savings over the individually optimized case.

The nature of synergy observed in this case is anchored in the fact that performing demand reductions with the retail buildings required significantly more energy consumption. These results seem plausible when considering the characteristics of each building. The retail building has approximately 21% less internal thermal capacitance (per unit area), 23% higher average construction U-value, six times higher ACH, and 2.6 times greater internal gains (per unit area) than the large office building. All these factors combine to make passive thermal storage in the retail building much less effective and more costly.

Aside from the gains in total cost savings, which were mostly attributable to reductions in energy expense in this case, the observed energy savings were significant as well. Scenario "a," "b," and "c," portfolio optimizations were able to achieve the desired target demand while consuming 17.5%, 5%, and 1.6% less energy than the aggregated individual optimizations, respectively.

Table 7.8: Case 1 total percent savings summary.

	No FR			FR		
Scenario	Rs+Ls	(RL)p	Diff.	Rs+Ls	(RL)p	Diff.
a	-19.47%	-26.98%	7.51	-21.15%	-28.44%	7.29
b	-24.76%	-26.98%	2.22	-26.20%	-28.44%	2.24
c	-26.54%	-26.98%	0.44	-27.83%	-28.44%	0.60

Chapter 8

Portfolio Case 2: Medium and Large Office Buildings

A second portfolio optimization case study was performed using medium office and large office buildings. The initial portfolio simulation study determined that the medium office and large office buildings were capable of achieving maximum demand reductions of 39.3 kW and 808 kW, respectively. Twenty medium office buildings were combined with a single large office building to create a portfolio where each building type has the ability shift equal amounts of load (i.e. approximately 800 kW). The twenty medium office buildings were assumed to be identical, and were included in the optimization through the use of a building multiplier to reduce problem dimensionality. Similar to Case 1, the portfolio-level target demand limit was set to encourage demand reductions without forcing all buildings into maximum demand reducing strategies. Figure 8.1 highlights the relationship between the portfolio NSU peak demand of 10 678 kW, the portfolio TDL of 9889 kW, and the minimum achievable peak of 9083 kW.

Figure 8.1 also shows that the portfolio-level demand reduction can be met entirely through load shifting via the medium office building, entirely via the large office building, or through some combination of shared load shedding. Similar to Case 1, three scenarios were investigated to evaluate under which conditions synergistic effect might be observed. Scenario "a" represents the case where the medium office building receives a low target demand limit such that it would perform the entire load reduction desired by the portfolio. In scenario "a" the large office building receives a high target demand limit such that reducing peak demand is not of primary concern. Scenario "b" represents the case where the medium office and large office buildings split the desired

portfolio demand reduction equally. Scenario "c" represents the case where the large office TDL is set such that it would be responsible for the entire portfolio demand reduction, while the medium office building receives a high target demand and need not reduce load. The target demand limits used to create these scenarios are shown in Table 8.1.

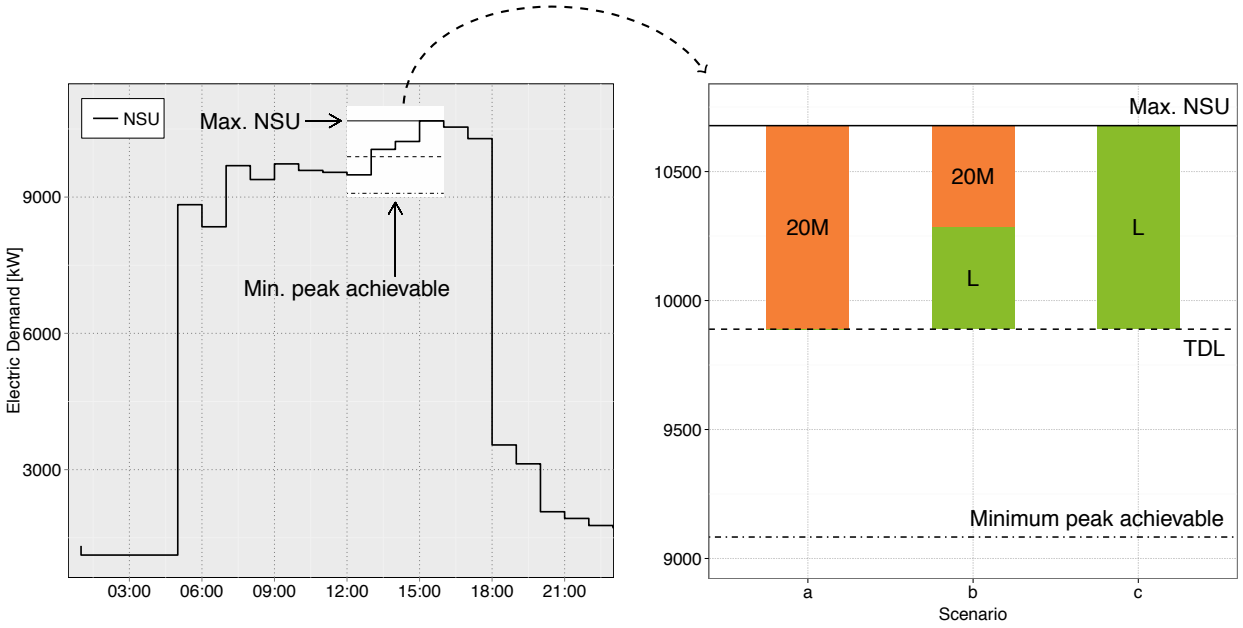

Figure 8.1: Case 2 scenarios and portfolio demand limits.

Table 8.1: Case 2 scenario target demand limits.

	Target Demand Limit			
Scenario	Ms	Ls	(ML)p	Units
a	6876	3013	9889	kW
b	7280	2609	9889	kW
c	7660	2229	9889	kW

8.1 Case 2a: Medium Office Requires Large Demand Reduction

As previously noted, scenario "a" represents the situation where the desired portfolio demand reduction can be entirely attributed to the peak management needs of the medium office building. This scenario could potentially arise if the large office building has previously set its peak demand and is now only concerned with maintaining operations below the peak. It seems possible that co-located buildings could peak on different days due to differences in operational patterns (e.g. one building is conditioned seven days a week while the other only five). To simulate this scenario the medium office building target demand limit was set at 6876 kW, requiring a load reduction of approximately 789 kW to avoid incurring a demand penalty. The target demand limit for the individual large office optimization was set at 3013 kW, which is nearly equivalent to the NSU peak, resulting in no demand penalty for NSU operation.

8.1.1 No Frequency Regulation

Figure 8.2 shows the medium office optimization results for scenario "a." The OPT Ms case began cooling at 1:00 AM to a minimum zone temperature of 22.2 °C at 6:00 AM. The temperature was slowly raised until the end of the peak period, achieving the desired demand reductions. Despite consuming 8822 kWh more energy than the NSU Ms case, $38 in energy cost savings were generated through shifting consumption to lower priced hours. In the portfolio optimization (OPT Mp), the medium office building began cooling one hour later than the OPT Ms case and in general kept zone temperatures higher. This ultimately resulted in less load reduction by the medium office building in the portfolio optimization.

The large office optimization results for scenario "a" are compared in Figure 8.3. The OPT Ls case started conditioning at hour ending 1:00 AM to take advantage of low energy prices and economizing opportunities. The building was cooled to a minimum of 23 °C at 7:00 AM and gradually increased to the upper bound of 23.9 °C at hour ending 3:00 PM. The light precooling strategy reduced demand well below the target level and generated $265 in energy cost savings through

energy arbitrage and a consumption reduction of 2326 kWh. A similar strategy was observed for the OPT Lp solution, however the portfolio optimization performed slightly more precooling reaching a minimum zone temperature of 22.5 °C at 6:00 AM.

The aggregate energy consumption, percent savings, and optimizer cost are summarized for scenario "a" in Figure 8.4. It is noted that the NSU cases are identical in energy consumption and total cost, implying that there is no natural synergy under NSU operation for this portfolio. This is due to the fact that the peak demand for both buildings under NSU control occurs during the same hour (i.e. no diversity). The top panel shows that the (ML)p solution was able to meet the demand target while consuming 4042 kWh (2.6%) less than the OPT Ms+Ls case. Since the target demand limit was set at the NSU peak, precooling was purely incentivized by real-time pricing signals and opportunities for arbitrage. Responding to price signals, however, resulted in a voluntary demand reduction by the large office building. In the portfolio optimization the medium office can "see" the beneficial opportunities for arbitrage in the large office building and advantageously utilize the voluntary demand reductions. The medium office is then required to perform less load reduction and consumes less energy. The portfolio solution resulted in $75 less total cost than the individual optimizations, which translates into a 0.7 percentage point increase in percent savings. Cost function values are presented in Table 8.2 for further reference.

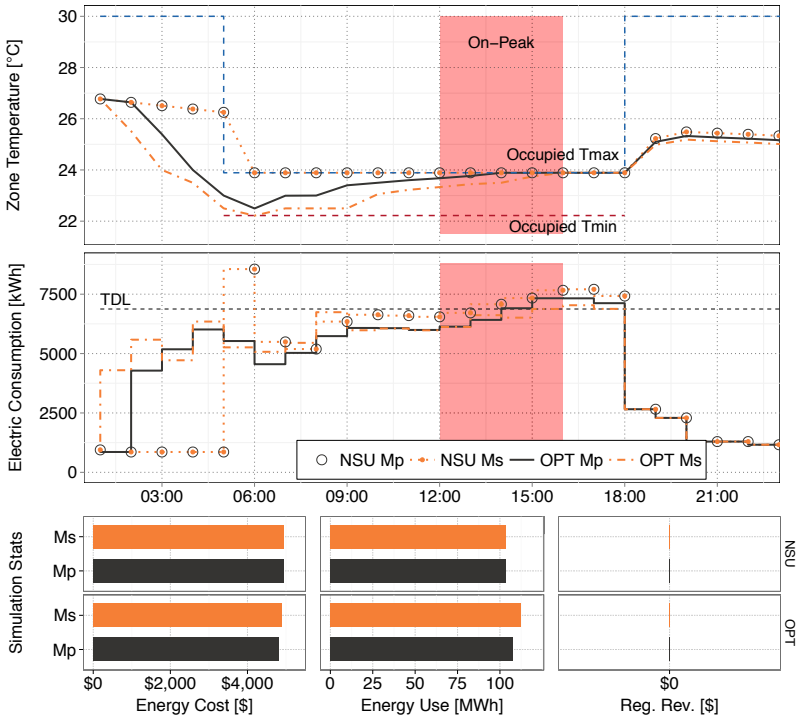

Figure 8.2: Case 2a medium office (x20) optimization results.

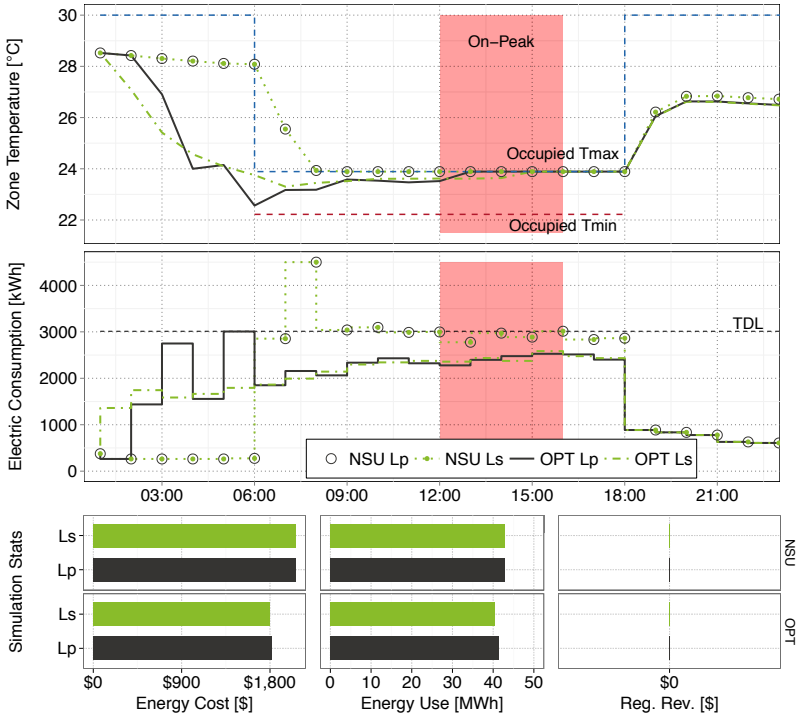

Figure 8.3: Case 2a large office (x1) optimization results.

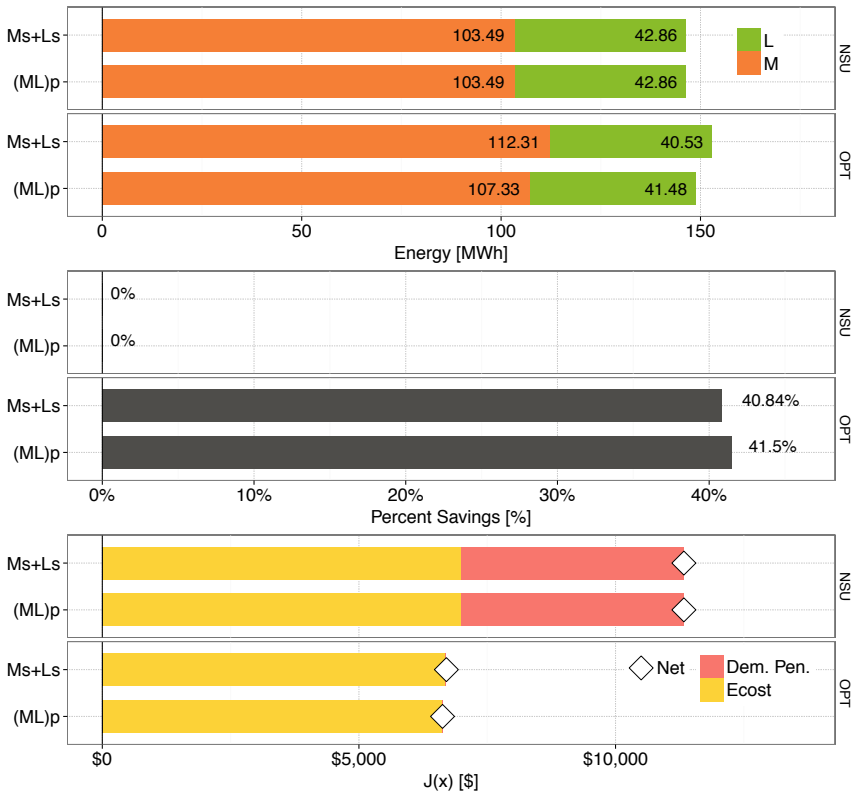

Figure 8.4: Case 2a aggregated individual and portfolio optimization results.

Table 8.2: Case 2a aggregated individual and portfolio optimization results.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	6992	6689	6992	6629
Energy Use [kWh]	146 352	152 848	146 352	148 806
Demand Penalty [$]	4339.75	15.28	4339.75	0
$J(\vec{x})$ [$]	11 331	6704	11 331	6629
% Difference	0%	-40.84%	0%	-41.50%

8.1.2 With Frequency Regulation

Scenario "a" optimizations were repeated with the inclusion of frequency regulation to evaluate the impact of ancillary service revenue on synergistic effect. The medium office and large office were assumed to be available to provide FR from 9:00 AM to 6:00 PM and 9:00 AM to 5:00 PM, respectively. The medium office optimization results are shown in Figure 8.5. The OPT Ms case performed slight precooling and then maintained the zone temperature at 23 °C during hours ending 9:00 AM to 11:00 AM. Since the medium office building was configured to provide FR via zone setpoint modulation, maintaining the temperature near the middle of the bounds allows for maximum regulation during these hours. The setpoint was lowered to 22.5 °C for hour ending 12:00 NOON to store additional cooling energy before the peak period. During the demand limiting period the temperature was raised towards the upper bound to achieve the desired demand reduction. The OPT Ms case was able to generate $269 of regulation revenue, while meeting the demand targets and only increasing energy expense by $44 over the NSU Ms case. No regulation is possible in the NSU cases since the upper temperature bound is maintained and the requirement of symmetric regulation was enforced.

A similar strategy was observed for the OPT Mp case, however conditioning did not begin until hour ending 4:00 AM and less precooling was performed overall. The optimizer kept the zone setpoint at 23 °C for an extra hour in order to achieve more regulation revenue. In total, the OPT Mp case generated $433 of regulation revenue.

The scenario "a" results for the large office building are provided in Figure 8.6. The NSU Ls and NSU Lp cases were able to generate $206 and $258 in regulation revenue, respectively. The portfolio NSU case was able to generate more regulation revenue due to the diversity created through aggregation and the fact that the medium office NSU does not provide FR. The OPT Ls case chose an early start-up with mild precooling before the on-peak period. Precooling was, once again, purely incentivized by reducing consumption during high-priced afternoon hours, since the TDL was set at the NSU peak. Voluntary demand reductions of 497 kW were observed as a

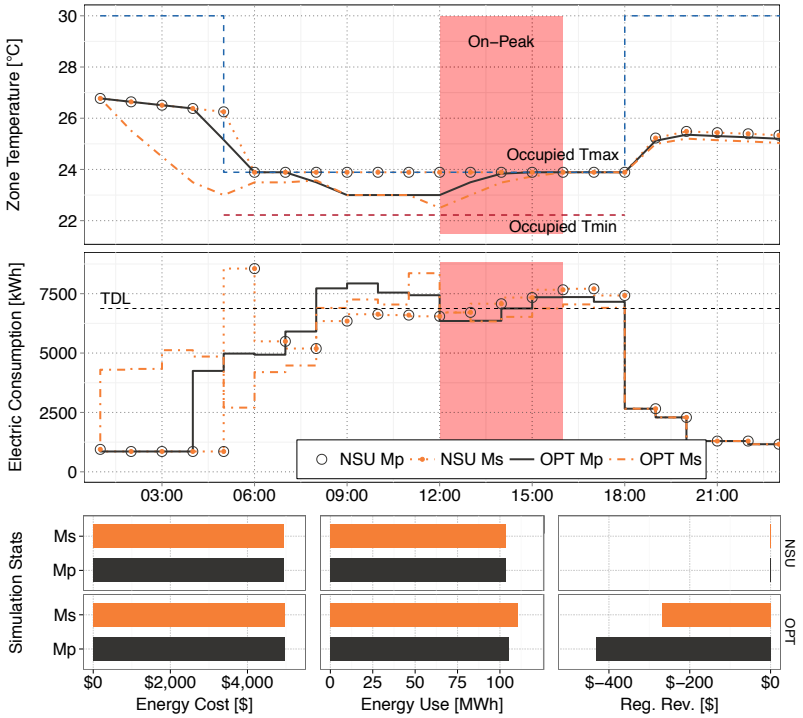

Figure 8.5: Case 2a medium office (x20) optimization results including FR.

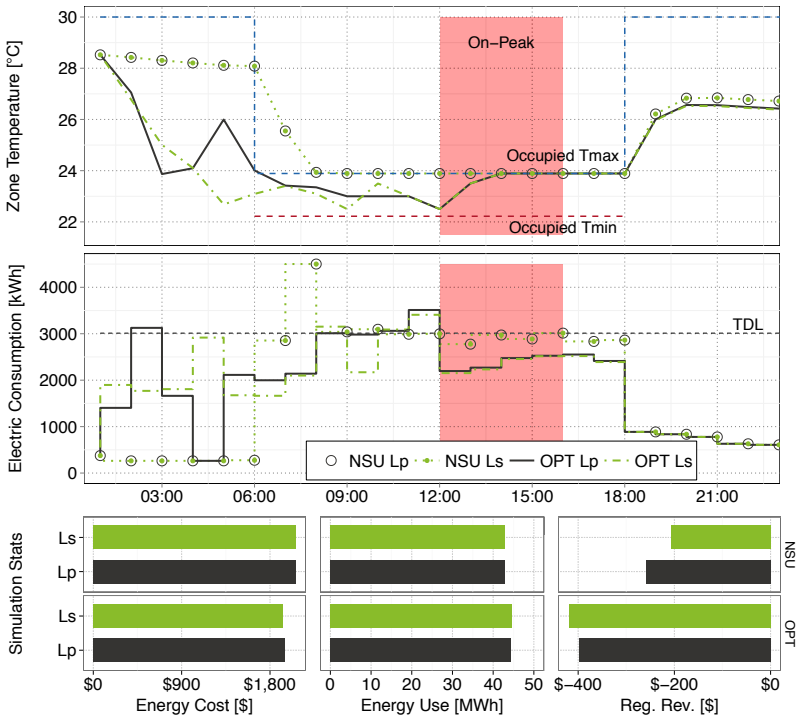

Figure 8.6: Case 2a large office (x1) optimization results including FR.

result of the large office building responding to hourly energy and regulation prices. The OPT Ls case was also able generate $418 in regulation revenue. The OPT Lp case chose a similar zone setpoint strategy as the OPT Ls case, although it resulted in slightly higher energy expense ($27) and slightly less regulation revenue ($20).

The aggregate optimization results for scenario "a" with frequency regulation are shown in Figure 8.7. The NSU cases shows that 1.81% to 2.28% savings were achieved by including FR in the NSU operation. Comparing the OPT cases shows that although the OPT Ms+Ls consumed more energy, the demand penalty was eliminated for nearly the same energy expense as the (ML)p case. The portfolio optimization spent only $20 (0.3%) more on energy, but generated $144 (21% more regulation revenue than the OPT Ms+Ls case. Further investigation of the simulation results showed that the OPT Mp setpoint strategy allowed for significantly more regulation potential in the medium office building from 9:00 AM to 12:00 NOON. This is presumably due to the HVAC system operating in a more flexible state (e.g. being further away from operating constraints). In total, the portfolio optimization resulted in an increase of 1.2 percentage points over the individual optimizations. Detailed cost function values are shown in Table 8.3.

8.2 Case 2b: Buildings Require Similar Demand Reductions

Scenario "b" represents the situation where individually the large office and medium office building contribute equally to the desired portfolio demand reduction (i.e. 400 kW each). This scenario may arise when both buildings are attempting to operate below their current monthly peak and avoid setting a new peak. This scenario was simulated by setting the Ms target demand at 7280 kW and the Ls target demand at 2609 kW. The portfolio TDL remained at 9889 kW thus portfolio optimization solutions are identical to scenario "a".

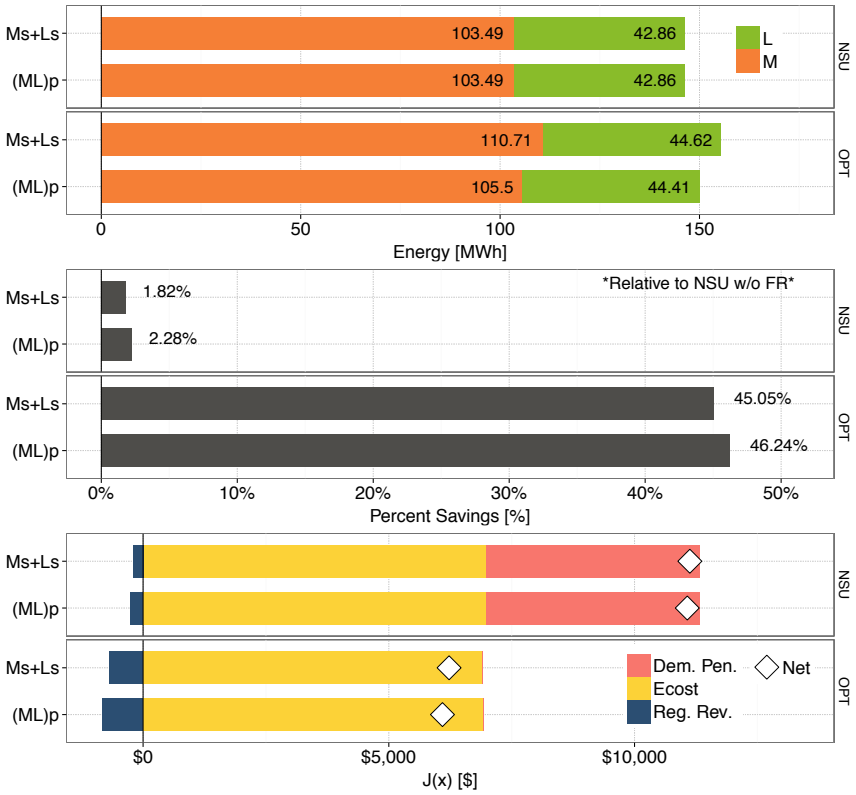

Figure 8.7: Case 2a aggregated individual and portfolio optimization results with FR.

Table 8.3: Case 2a aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	686.65	258.19	830.57
Energy Cost [$]	6992	6901	6992	6922
Energy Use [kWh]	146 352	155 327	146 352	149 914
Demand Penalty [$]	4339.75	12.26	4339.75	0
$J(\vec{x})$ [$]	11 125	6227	11 073	6091
% Difference	-1.82%	-45.05%	-2.28%	-46.24%

8.2.1 No Frequency Regulation

The results for the scenario "b" medium office building optimizations are shown in Figure 8.8. Mild precooling was observed in the OPT Ms case with the zone setpoint held near the middle of the occupied temperature bounds until hour ending 8:00 AM. The zone temperature was slowly raised towards the upper bound reducing the NSU peak demand by 386 kW. Aside from peak reductions, $120 in energy cost savings over the NSU case were achieved through arbitrage, despite higher energy consumption. The temperature and demand profiles appear nearly identical between the OPT Ms and OPT Mp cases. Further investigation of the simulation details shows a difference in energy consumption of 109 kWh, which translates to only $6 (0.1%) difference in energy expense.

The large office optimization results for scenario "b" are shown in Figure 8.9. The OPT Ls case performed slightly more precooling than in scenario "a" to meet the lower demand target. Thermal energy was stored by maintaing the zone setpoint near 23.3 °C preceding the on-peak period, and then released by raising to the upper bound during the peak period. The OPT Lp solution appears quite similar to the OPT Ls case. Further investigation of the simulation details shows that there is only a 1.6% (660 kWh) difference in energy consumption and 2% ($8) difference in energy cost.

Aggregate results for scenario "b" are shown in Figure 8.10. The NSU cases are identical in energy consumption and total cost implying that no natural diversity exists under NSU operation. Combining the results of the individual optimizations achieved nearly identical results as the portfolio optimization. The OPT (ML)p total cost was $15 higher than that of the aggregated individual optimizations, which translates into a 0.1 percentage point difference in total percent savings. Since the portfolio optimization has the ability to choose a solution identical to the individual optimizations, an optimal portfolio solution is expected to have a final cost equal to or lower than the aggregation of individual optimizations. The main differences between the portfolio and individual solutions occur during hours preceding 7:00 AM. The average energy price during these hours is $0.02/kWh making it difficult for the optimizer to discriminate between small variations in

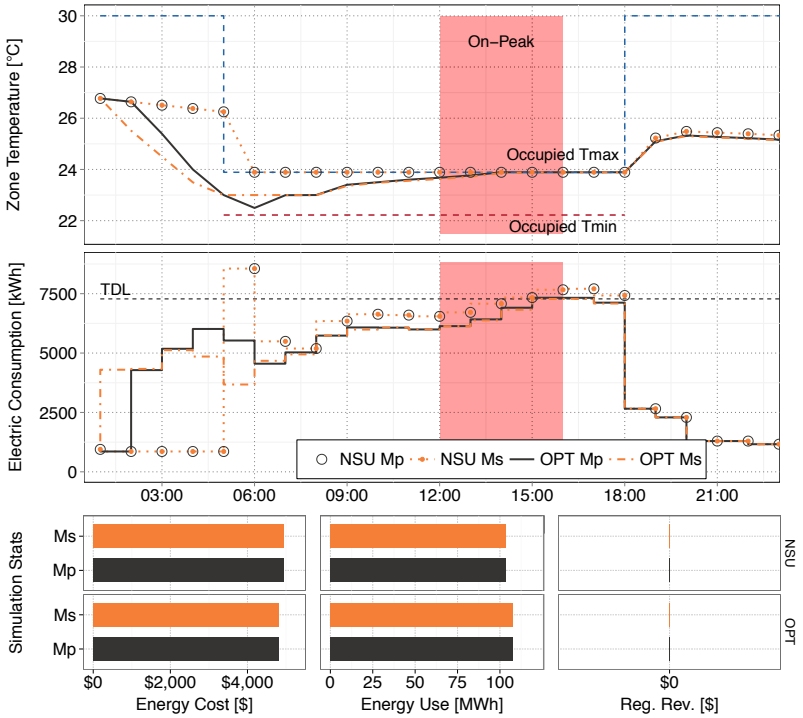

Figure 8.8: Case 2b medium office (x20) optimization results.

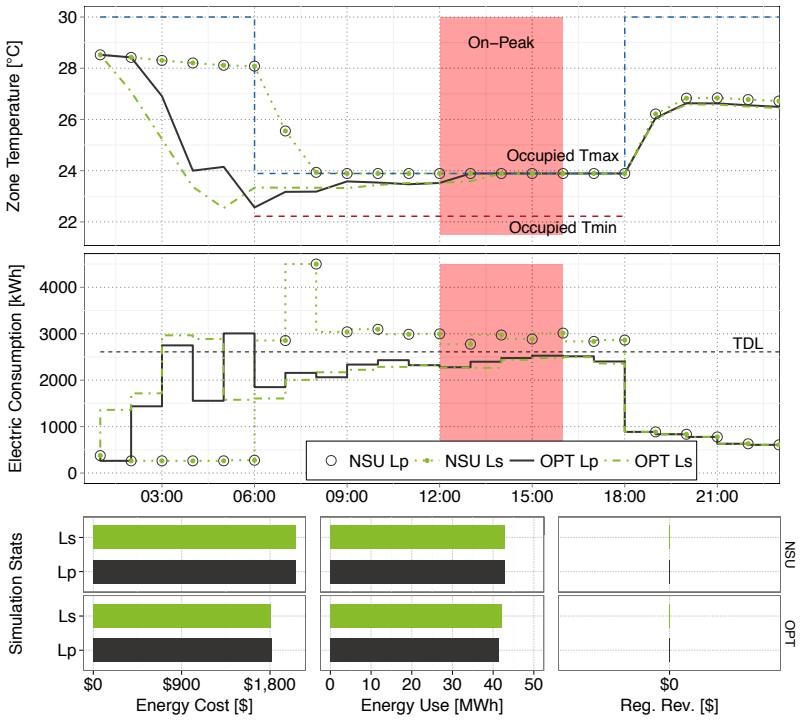

Figure 8.9: Case 2b large office (x1) optimization results.

consumption during these hours. The portfolio optimization may have been prematurely stopped on a slightly suboptimal solution. However, for practical purposes, the portfolio optimization very nearly reproduced the individual solutions of scenario "b." Cost function numbers are shown in Table 8.4 for reference.

8.2.2 With Frequency Regulation

Scenario "b" was repeated with the inclusion of frequency regulation to evaluate the impact of ancillary service revenue on the optimizations. The medium office optimizations are shown in Figure 8.11, and exhibit a somewhat different solution than that observed in the non-FR OPT Ms case. Near NSU operation was followed until hour ending 7:00 AM. For hours ending 9:00 to 12:00 NOON the zone setpoint was held near 23 °C in order to maximize regulation during these hours. The optimizer determined that potential regulation revenue outweighed the opportunity for energy expense savings in this scenario. Providing FR required consuming 2799 kWh more than the NSU Ms case, costing an additional $35. A similar strategy was observed for the OPT Mp solution. From the optimizer perspective, the sum of regulation revenue and energy cost differ by only $0.66 between the medium office individual and portfolio solutions. However, if not aggregated with the large office building, the OPT Mp solution would have exceeded the medium office TDL and incurred a demand penalty.

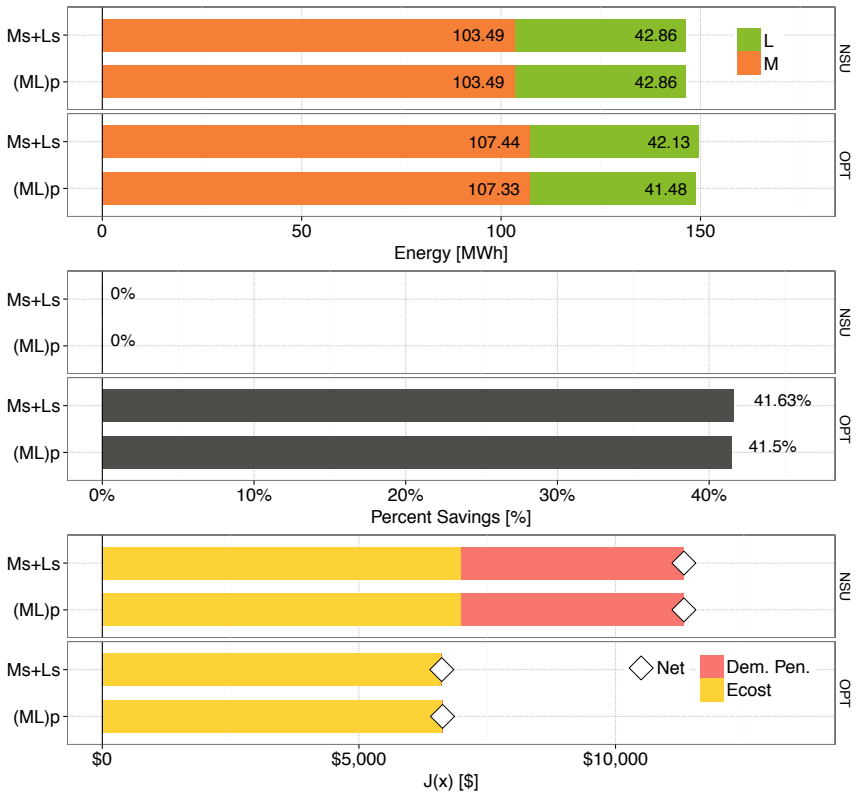

Figure 8.10: Case 2b aggregated individual and portfolio optimization results.

Table 8.4: Case 2b aggregated individual and portfolio optimization results.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	6992	6614	6992	6629
Energy Use [kWh]	146 352	149 575	146 352	148 806
Demand Penalty [$]	4339.75	0	4339.8	0
$J(\vec{x})$ [$]	11 331	6614	11 331	6629
% Difference	0%	-41.63%	0%	-41.50%

The large office results for scenario "b" including frequency regulation are shown in Figure 8.12. As in scenario "a", the NSU Ls and NSU Lp cases generated $206 and $258 in regulation, respectively. The OPT Ls case performed fairly significant precooling with zone temperatures held near the lower temperature bound for four hours preceding the peak period. The OPT Ls case generated $177 more regulation revenue than the NSU Ls case, and achieved $110 in energy cost savings despite consuming 2856 kWh more. The OPT Lp case, in general, precooled less keeping warmer zone temperatures throughout the morning. The OPT Lp case consumed 1303 kWh less than the OPT Ls case, and achieved slightly more regulation revenue (i.e. $14).

The aggregate simulation results are highlighted in Figure 8.13. Energy consumption, energy cost, and demand penalties were identical between the NSU cases. The additional regulation performed by the portfolio NSU case equated to an increase of 0.46 percentage points in percent savings. The OPT (ML)p case was able to achieve the target demand by consuming about 1.5% less energy, while generating an additional $16 in regulation revenue. Overall the savings only represent an additional 0.1 percentage point increase in total percent savings. For practical purposes the portfolio and individual optimization results produce nearly identical results. The cost function details are highlighted Table 8.5.

8.3 Case 2c: Large Office Requires Large Demand Reduction

Scenario "c" represents the situation where individually the large office desires to reduce demand by approximately 784 kW, while the medium office building has no need to reduce load. This scenario may arise when the medium office building has previously set its monthly peak, while the large office building is experiencing a potentially peak-setting day. This situation could occur due to operational differences such as whether or not a building is conditioned seven days a week. This scenario was simulated by setting the large office TDL at 2229 kW and the medium office TDL at 7660 kW. The portfolio TDL remained at 9889 kW thus portfolio optimization results are identical to scenarios "a" and "b".

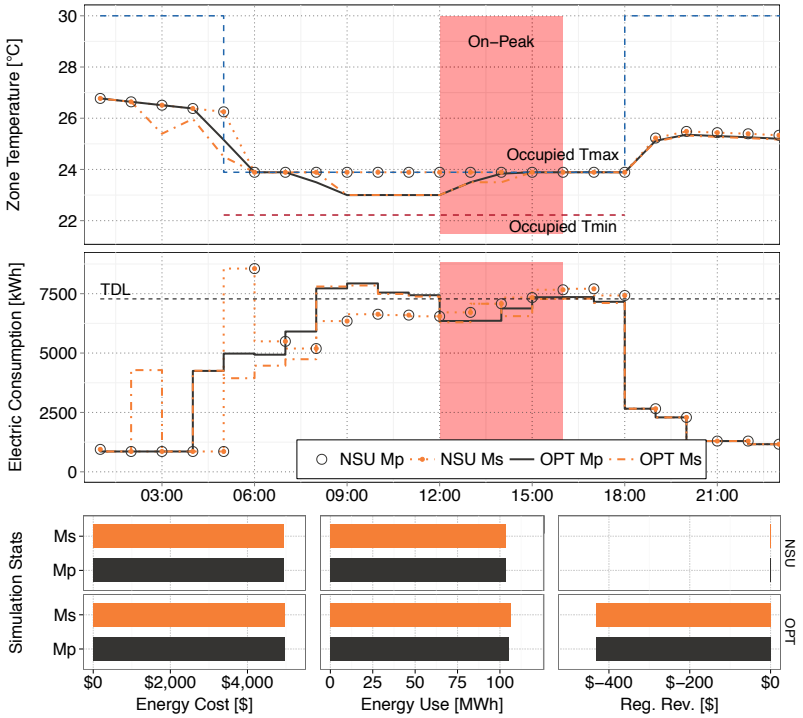

Figure 8.11: Case 2b medium office (x20) optimization results including FR.

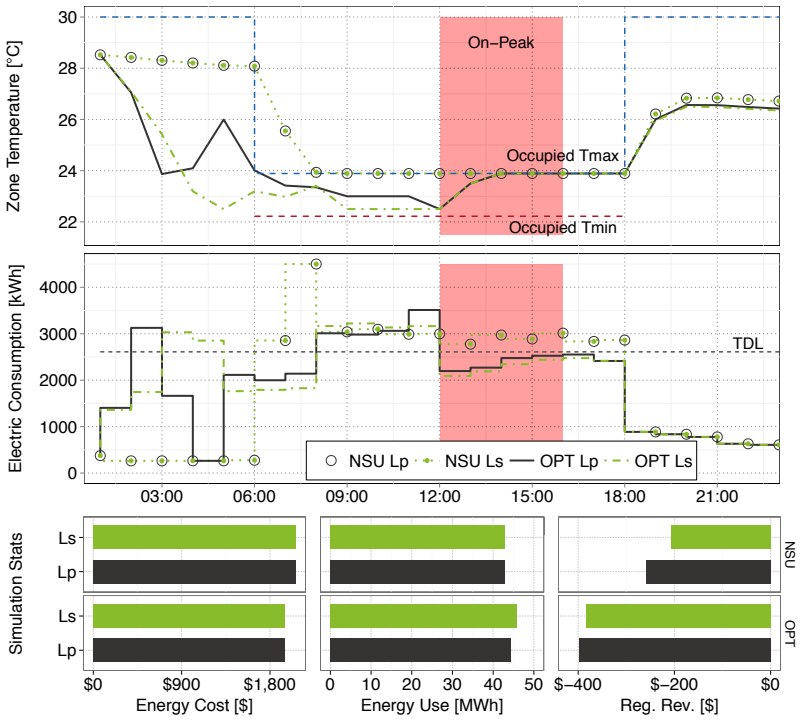

Figure 8.12: Case 2b large office (x1) optimization results including FR.

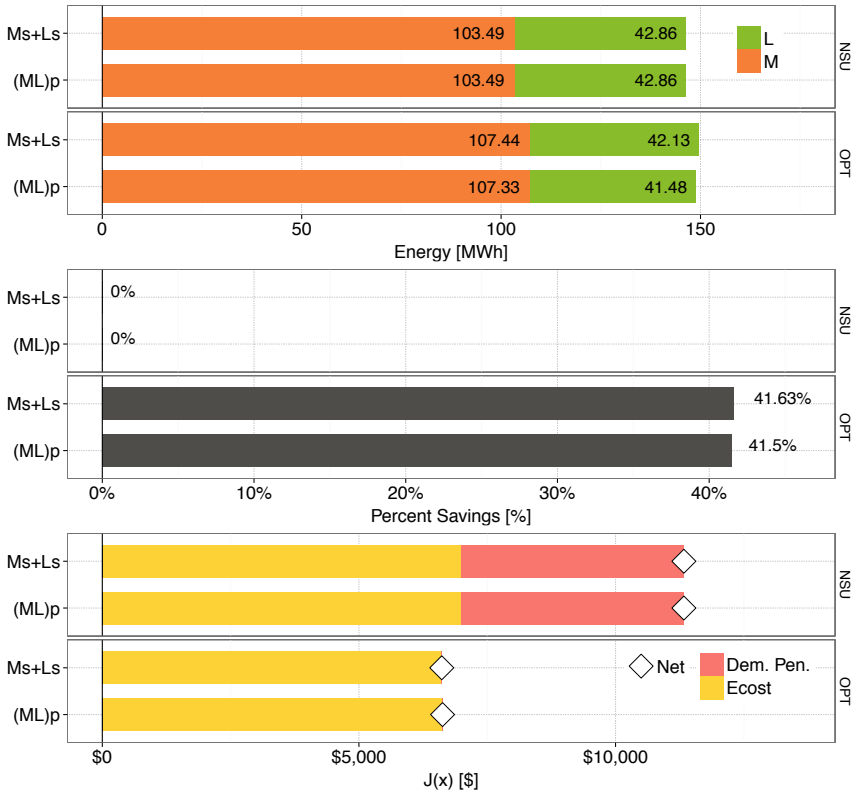

Figure 8.13: Case 2b aggregated individual and portfolio optimization results.

Table 8.5: Case 2b aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	814.56	258.19	830.57
Energy Cost [$]	6992	6916	6992	6922
Energy Use [kWh]	146 352	152 007	146 352	149 914
Demand Penalty [$]	4339.75	0	4339.8	0
$J(\vec{x})$ [$]	11 125	6102	11 073	6091
% Difference	-1.82%	-46.15%	-2.28%	-46.24%

8.3.1 No Frequency Regulation

The medium office optimization results for scenario "c" are shown in Figure 8.14. The OPT Ms case chose to start conditioning the space early and slightly precool during the first few hours of occupancy. This resulted in the consumption of 506 kWh more than the NSU Ms case with the benefit of saving $133 in energy cost. In comparison, the OPT Mp performed more precooling, achieving greater demand reductions at higher energy consumption and cost.

The large office building scenario "c" optimization results are highlighted in Figure 8.15. Due to the low target demand, the OPT Ls case significantly precools the zone and maintains the lower temperature bound of 22.2 °C preceding the on-peak period to achieve the required load reduction. The load shedding requires 4586 kWh more than the NSU Ls case, costing an additional $133. In comparison, the portfolio optimization chose to perform less load shifting with the large office building, keeping zone temperatures much higher during the morning hours.

The aggregated results are summarized in Figure 8.16. The NSU cases were identical, thus, no diversity exists through aggregation under NSU control. The OPT (ML)p case was able to eliminate the demand penalty while consuming 2638 kWh (1.75%) less energy than the OPT Ms+Ls case, saving an additional $96 (1.4%) in energy cost. Overall the portfolio optimization resulted in a 0.9 percentage point increase in percent savings over the individually optimized case. The cost function numbers are presented in Table 8.6 for reference.

8.3.2 With Frequency Regulation

Scenario "c" results were also repeated with the inclusion of frequency regulation in the optimization. Figure 8.17 highlights the optimization results for the medium office building including frequency regulation. The scenario "c" strategy was nearly identical to that observed for scenario "b," with the exception of a one hour later start-up in the present example. The zone temperatures were held near 23 °C from 9:00 AM to 12:00 NOON in order to provide maximum regulation preceding the peak period. The higher demand limit in scenario "c" allowed the OPT Ms case to

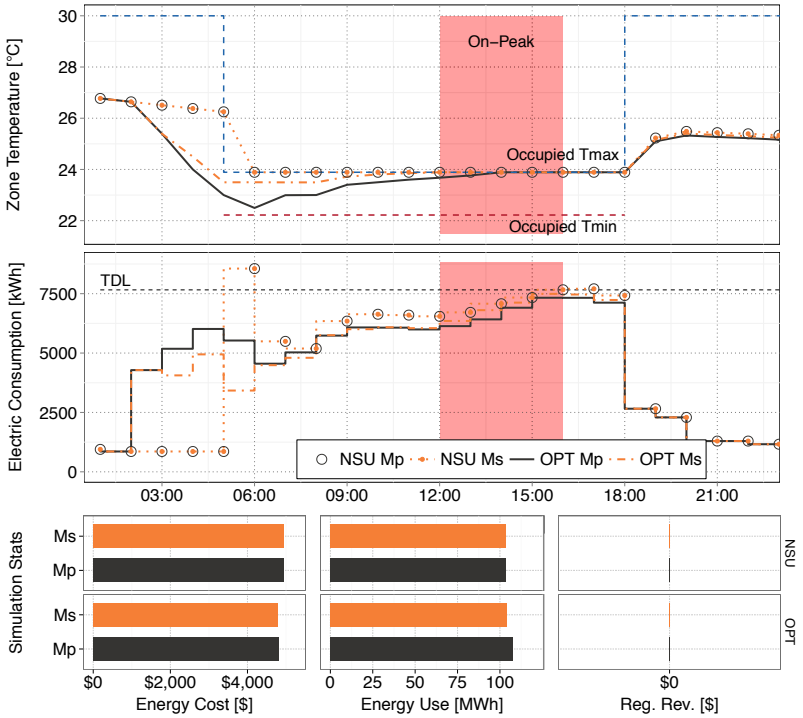

Figure 8.14: Case 2c medium office (x20) optimization results.

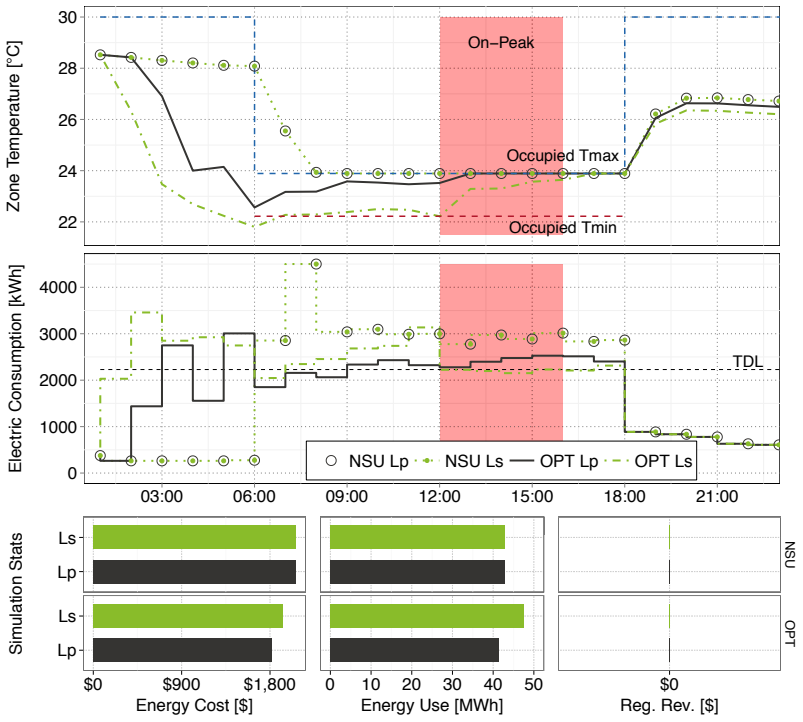

Figure 8.15: Case 2c large office (x1) optimization results.

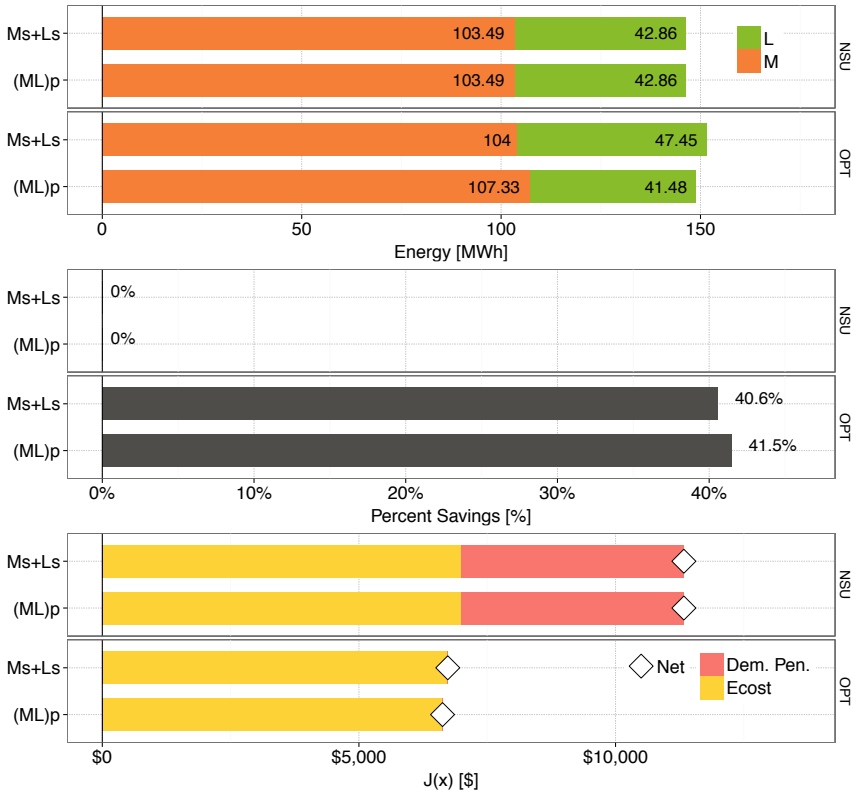

Figure 8.16: Case 2c aggregated individual and portfolio optimization results.

Table 8.6: Case 2c aggregated individual and portfolio optimization results.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	6992	6725	6992	6629
Energy Use [kWh]	146 352	151 443	146 352	148 806
Demand Penalty [$]	4339.75	6.06	4339.75	0
$J(\vec{x})$ [$]	11 331	6731	11 331	6629
% Difference	0%	-40.60%	0%	-41.50%

generate $44 more revenue than the scenario “b” case.

The large office results for scenario “c” are presented in Figure 8.18. As in the non-FR example, the low target demand limit forces extensive precooling in order to reduce on-peak demand in the OPT Ls case. Although load reduction was the primary objective, $271 in regulation revenue was still generated. In comparison the OPT Lp case performed much less precooling, consuming less energy and providing almost $400 in regulation revenue.

Total energy consumption, percent savings, and optimizer cost are shown in Figure 8.19 to compare the aggregate results. The portfolio optimization (i.e. OPT (ML)p) was able to eliminate the demand penalty while consuming 4384 kWh (3%) less energy. The portfolio optimization also resulted in $24 less energy cost and $84 (11%) more revenue. Further investigation of simulation results shows that the adhering to the low target demand in the OPT Ls case significantly reduces its FR potential during hours ending 1:00 PM to 4:00 PM, when compared to the portfolio optimized solution. Overall the portfolio optimization resulted in a percent savings 1.0 percentage point higher than the individually optimized case. Cost function values are summarized in Table 8.7 for additional reference.

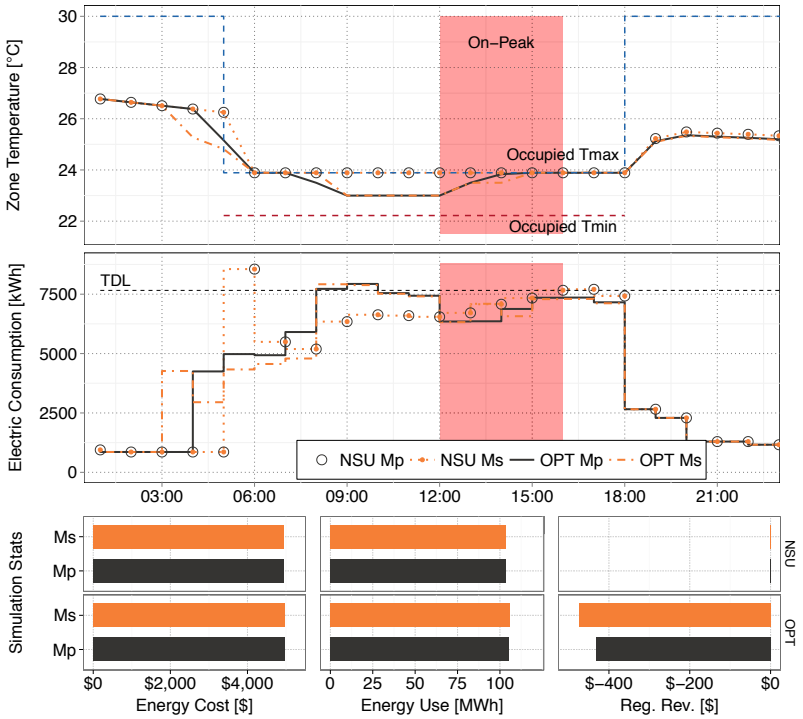

Figure 8.17: Case 2c medium office (x20) optimization results including FR.

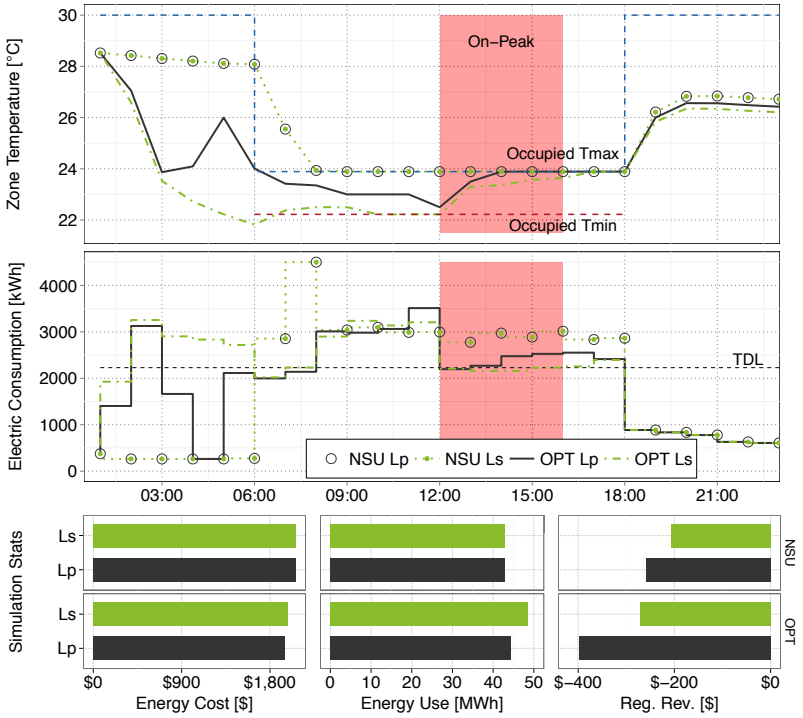

Figure 8.18: Case 2c large office (x1) optimization results including FR.

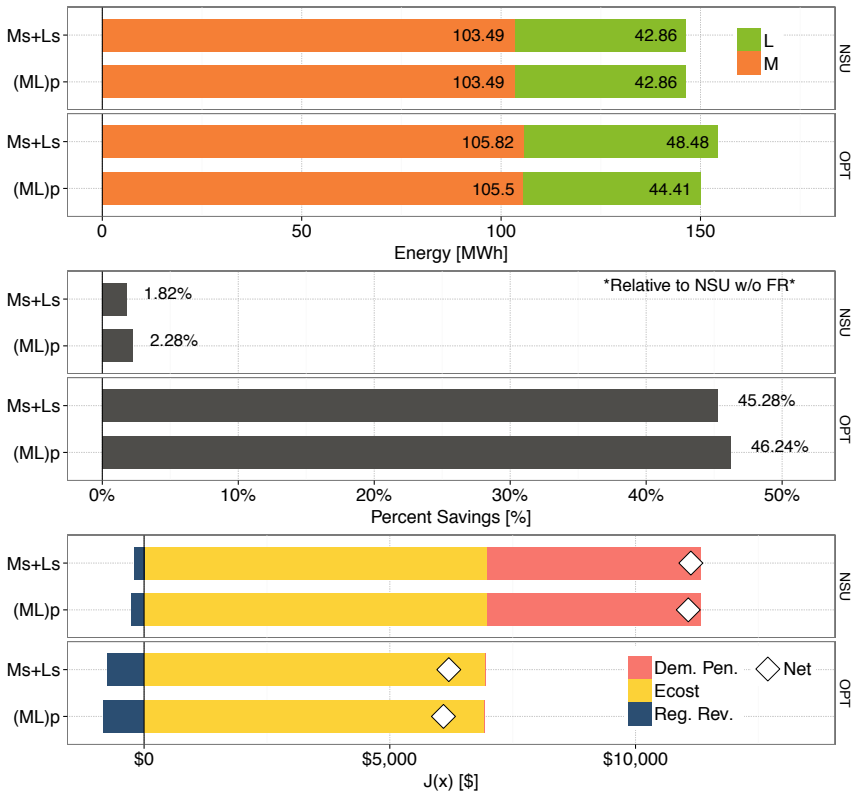

Figure 8.19: Case 2c aggregated individual and portfolio optimization results with FR.

Table 8.7: Case 2c aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	746.15	258.19	830.57
Energy Cost [$]	6992	6947	6992	6922
Energy Use [kWh]	146 352	154 299	146 352	149 914
Demand Penalty [$]	4339.75	0	4339.75	0
$J(\vec{x})$ [$]	11 125	6201	11 073	6091
% Difference	-1.82%	-45.28%	-2.28%	-46.24%

8.4 Case 2 Summary

Comparing the results across scenarios "a", "b", and "c" for Case 2, again highlights that the level of synergy may be somewhat dependent on the conditions of the individual building optimizations before aggregation. Table 8.8 summarizes the achieved precent savings by the OPT cases for the investigated scenarios. The NSU Ms+Ls case was identical for all scenarios, which allows for fair comparison between percent savings across scenarios. The portfolio results determined that for this specific portfolio construction (i.e. 20 medium office, and one large office) and simulation case the optimal operation was achieved when the load shift was split between the buildings. Scenario "b" virtually reproduced the portfolio optimized results showing a difference of only 0.1 percentage points in total percent savings. If the individual buildings had been under conditions such as scenario "a" or "c," slight synergy would exist among the portfolio.

Table 8.8: Case 2 total percent savings summary.

	No FR			FR		
Scenario	Ms+Ls	(ML)p	Diff.	Ms+Ls	(ML)p	Diff.
a	-40.84%	-41.50%	0.66	-45.05%	-46.24%	1.20
b	-41.63%	-41.50%	-0.13	-46.15%	-46.24%	0.09
c	-40.60%	-41.50%	0.90	-45.28%	-46.24%	0.96

The observed synergies in this example tend to hinge on the fact that both buildings voluntarily reduce demand in response to hourly prices (and perhaps regulation prices) when given a high target demand (i.e. no demand penalties). In scenario "a," the large office building would naturally reduce peak demand by 428 kW as a result of taking advantage of price arbitrage. Similarly, in scenario "c" the medium office building would reduce peak demand by 172 kW. When optimized individually, these voluntary load reductions remain undisclosed to the other buildings and cause greater load reduction than necessary. Put another way, the centralized optimization allows the voluntary demand reductions to be appropriated to the desired demand reduction, eliminating

excess load shifting. Considering the total voluntary demand reductions available between both buildings gives an initial reduction of 600 kW, or 76% of the desired reduction. Only an additional 189 kW reduction needed to be encouraged via demand penalty. By making use of all voluntary load reductions the portfolio optimizations in this example were able to achieve the desired TDL at lower energy consumption than the aggregated individual optimizations.

The optimizer chose to allocate 131 kW of the remaining 189 kW to the medium office buildings and 58 kW to the large office building in the non-FR examples. The fact that the remaining reductions were not allocated to a single building may suggest that it was more efficient to split the workload load, in this case, since the buildings are somewhat similar.

Aside from energy savings from the portfolio optimization, regulation revenue was also higher for scenarios consider in this case. Striving to achieve large individual demand reductions tended to keep equipment or setpoints closer to operating constraints, resulting in less ability to provide frequency regulation.

Although the potential for interesting synergy was observed, the additional savings were somewhat small. Reductions in energy and increases in FR typically only resulted in approximately one percentage point increase in the total portfolio percent savings. This is likely due to the fact that the additional energy consumption often occurred during low priced hours and that regulation revenue was typically only one tenth (or less) of the magnitude of the energy cost. The observed 3% reduction in end-use electricity consumption, however, may be considered significant.

Chapter 9

Portfolio Case 3: Medium and Large Office Buildings with High Demand Limit

In the simulation case studies presented in Chapters 7 and 8, the portfolios were required to significantly reduce on-peak demand. When demand reductions are the primary objective, frequency regulation becomes secondary due to the magnitude of the demand penalties applied. This nature was observed in the relatively minimal impact of including frequency regulation in Cases 1 and 2. It was desired to further investigate scenarios when ancillary service revenues may lead to greater benefit and actively participate in synergistic effect. To this end, the portfolio considered in Chapter 8, consisting of twenty medium office buildings and one large office building, was used to perform a third portfolio simulation study.

Case 3 primarily differs from previous cases in that a higher portfolio target demand was applied so that only frequency regulation excursions could potentially increase peak demand. Previous simulation results (e.g. the high TDL results of Chapter 6) showed that when peak demand is of no concern the optimization often seeks to maximize regulation opportunities. In the new simulation case, the TDL was set to constrain the portfolio from achieving the maximum possible FR solution while having no impact on operations without FR. It seems conceivable that this configuration may create opportunities for buildings to voluntarily reduce load in order to increase FR capability and simultaneously adhere to peak demand limits. Furthermore, opportunities may also exist for buildings to provide differing amounts of load reduction and frequency regulation in pursuit of a greater portfolio solution (i.e. synergy).

The portfolio target demand limit is shown in relation to demand profiles and regulation peaks

in Figure 9.1. The NSU demand profile represents the baseline power draw for the portfolio under NSU operation, while the Max NSU+FR threshold represents the maximum possible power draw when providing maximum regulation under NSU control. The OPT (No TDL) profile represents the optimal baseline when no demand penalty is employed which, allows the buildings to generate the maximum possible regulation revenue. The Max OPT+FR threshold represents the maximum power draw when the OPT case is providing regulation.

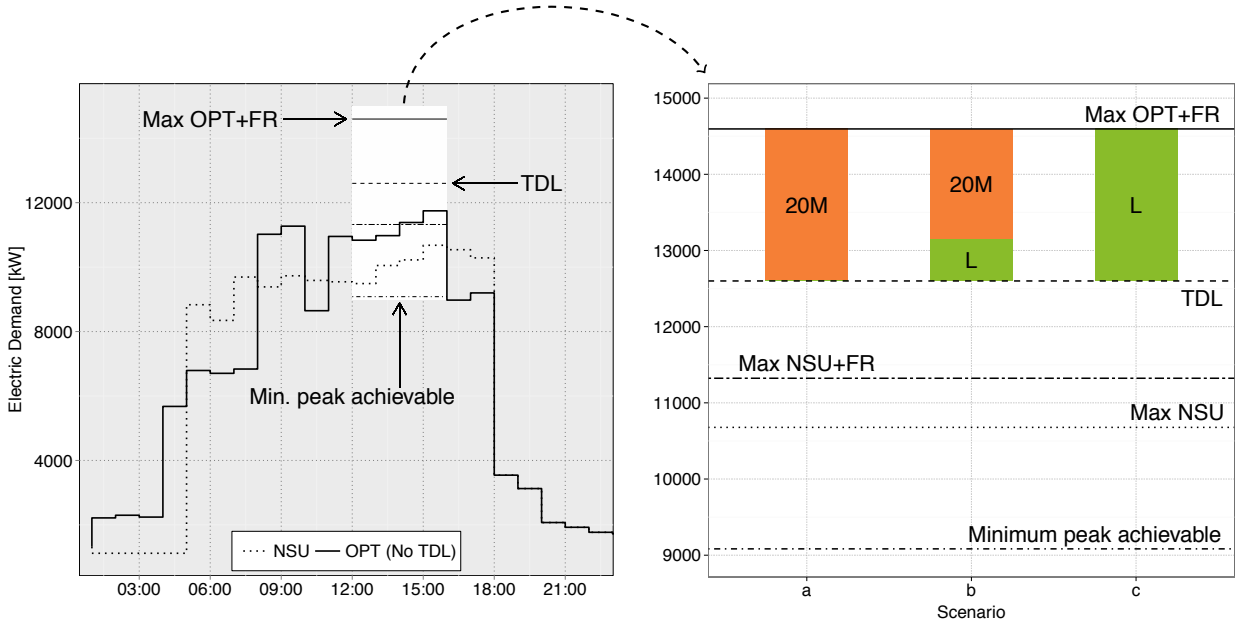

Figure 9.1: Case 3 scenarios and portfolio demand limits.

The portfolio TDL was set at 12 600 kW, which is below the 14 467 kW OPT+FR maximum and above the 11 323 kW NSU+FR maximum. This TDL allows NSU operation to achieve maximum regulation without penalty, but prohibits the maximum FR strategy (i.e. when no limits are applied). With respect to the maximum FR solution, the portfolio optimization must determine an optimal strategy with an on-peak demand 1867 kW lower than the max FR solution. These demand reductions can be achieved through a variety of reductions in FR and altering the baseline power draw through load shifting.

Three scenarios were investigated to explore synergistic effect in relation to different individual building conditions. Scenario "a" represents the situation where the large office building

Table 9.1: Case 3 scenario target demand limits.

	Target Demand Limit			
Scenario	Ms	Ls	(ML)p	Units
a	8337	4263	12 600	kW
b	8880	3720	12 600	kW
c	10 340	2260	12 600	kW

is allowed to attain its maximum regulation solution, while the medium office is constrained to a lower TDL. Scenario "c" represents the situation where the medium office can reach its maximum regulation strategy, while the large office building provides peak demand relief. Scenario "b" represents the situation where neither building can achieve their respective maximum FR strategies without violating personal TDL's. The asymmetric distribution was chosen to allow the NSU cases to still provide maximum FR without incurring demand penalties. The target demand limits for each scenario are shown in Table 9.1.

9.1 Case 3a: Medium Office Peak Constrained

As earlier noted, scenario "a" represents the situation when the medium office has a peak demand constraint, while the large office building is able to attain its maximum regulation strategy. This situation may arise if the large office building has previously set a peak much higher than the current operating day while the medium office is operating closer to its peak. Variations in operational patterns and control strategies throughout the month could conceivably produce such a scenario. This scenario was simulated by applying target demand limits of 8337 kW and 4263 kW to the medium and large office buildings, respectively.

9.1.1 No Frequency Regulation

Figure 9.2 highlights the optimization results for the medium office building. The OPT Ms and OPT Mp solutions started conditioning the building three hours before the NSU case, slightly cooling below the upper temperature limit from hours ending 6:00 AM to 9:00 AM. Since regulation was not included and the TDL is higher than the NSU peak, non-NSU operations were purely incentivized by hourly energy prices. Energy arbitrage in the OPT Ms and OPT Mp cases was able to save $133 in energy expense over NSU operation, despite consuming 506 kWh more energy.

The large office optimization results for scenario "a" are shown in Figure 9.3. The OPT Ls case shows an early start with zone temperatures maintained slightly below the upper temperature bound. This resulted in decreased energy consumption during high priced hours and an overall energy expense savings of $265. A similar solution was observed for the OPT Lp case, however, more precooling occurred from 4:00 AM to 6:00 AM, resulting in slightly higher energy consumption for the portfolio strategy.

The aggregate energy consumption, percent savings, and optimizer cost for scenario "a" are summarized in Figure 9.4. The NSU cases were identical in energy consumption and optimizer cost. This implies that no natural diversity existed when portfolio members were operated under NSU control since the individual peaks occured coincident in time. Economizing opportunities and HVAC efficiency improvements allowed the OPT Ms+Ls case to consume 1820 kWh (1.2%) less energy than the NSU case. Additionally, more energy was consumed during lower priced hours resulting in $398 (5.7%) savings in energy expense.

The OPT (ML)p solution consumed slightly more energy than the aggregate individual optimizations, and consequently shows 0.2 percentage points less savings. Since the portfolio optimization has the opportunity to choose solutions equivalent to the individual optimizations, an optimal portfolio solution should always result in savings greater than or equal to the aggregate individual optimizations. In this case the portfolio optimization prematurely stopped on a slightly suboptimal solution. The main difference between the individual and portfolio solutions occurred in the

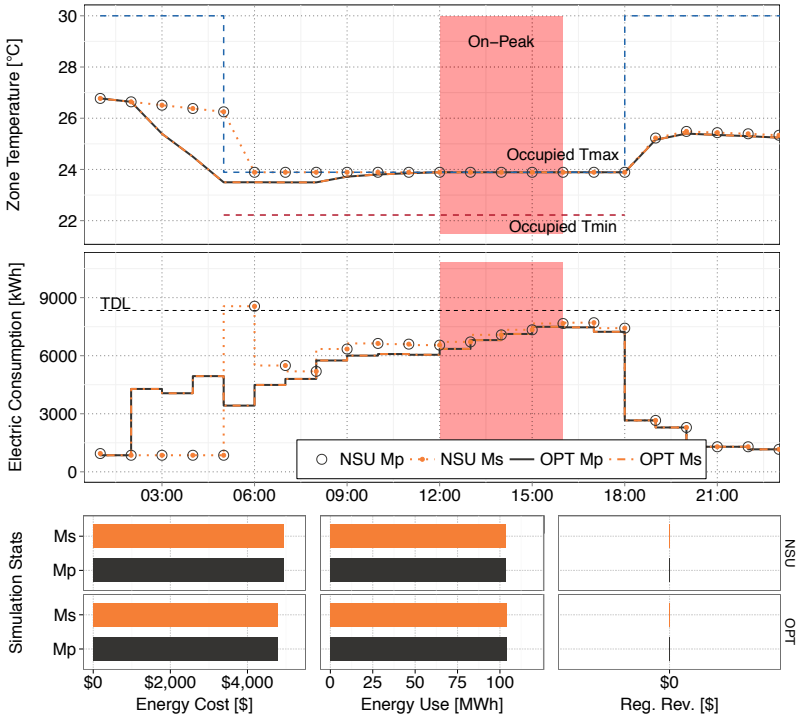

Figure 9.2: Case 3a medium office (x20) optimization results.

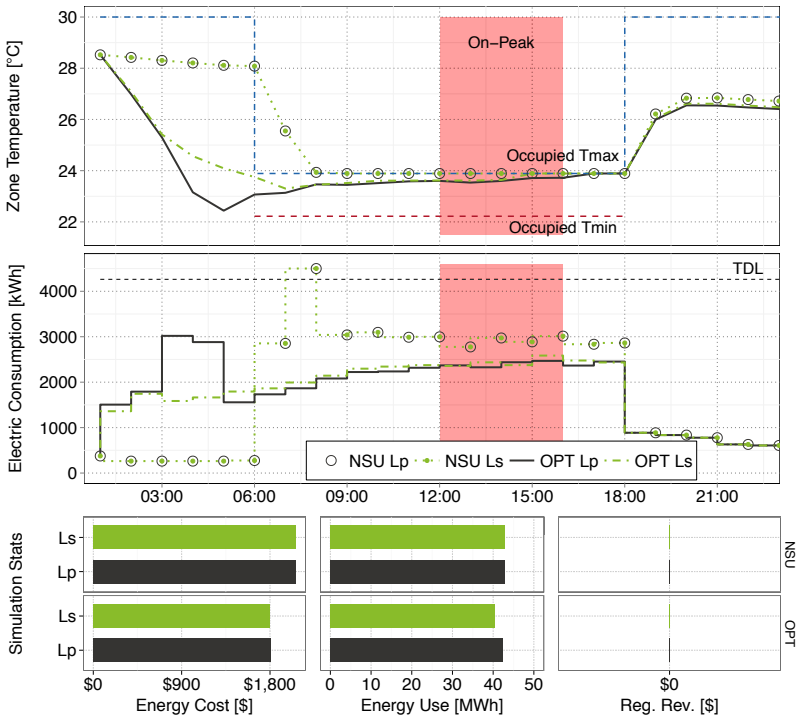

Figure 9.3: Case 3a large office (x1) optimization results.

Figure 9.4: Case 3a aggregated individual and portfolio optimization results.

Table 9.2: Case 3a aggregated individual and portfolio optimization results.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	6992	6593	6992	6605
Energy Use [kWh]	146 352	144 532	146 352	146 336
Demand Penalty [$]	0	0	0	0
$J(\vec{x})$ [$]	6992	6593	6992	6605
% Difference	0%	-5.70%	0%	-5.53%

large office building from 4:00 AM to 6:00 AM. The average energy price for these hours was only \$0.02/kWh making it difficult for the optimizer to discriminate between variations in consumption during this period. The difference in energy consumption resulted in only \$11 difference in energy expense. The cost function numbers are provided in Table 9.2 for reference.

9.1.2 With Frequency Regulation

Scenario "a" results were repeated with the inclusion of frequency regulation in the optimization. Results for the medium office building are highlighted in Figure 9.5. The OPT Ms case began cooling two hours before normal NSU operation and held the upper temperature boundary until hour ending 9:00 AM. At 9:00 AM, the building became available to provide FR and the setpoint was lowered to the middle of the temperature bounds in order to provide maximum regulation. A temperature of 23.5 °C was held during the on-peak period to strike a balance between FR capability and peak demand management. The OPT Ms case was able to generate \$650 in regulation revenue while consuming only 3161 kWh (3%) more energy than the NSU Ms case. The additional energy consumption resulted in an increase in energy expense of \$93 (1.9%). A similar solution was observed for the OPT Mp case, however zone temperatures were maintained at 23 °C for the entire on-peak period. This allowed for the generation of more regulation revenue, at the cost of slightly higher energy consumption and expense.

The results for the large office scenario "a" optimization including FR are shown in Figure 9.6. The individual TDL in this case was high enough to allow the OPT Ls case to achieve its maximum regulation strategy without incurring a demand penalty. The OPT Ls started conditioning at 1:00 AM and kept zone temperatures near the upper bound for the morning hours. From hours ending 1:00 PM to 4:00 PM the zone temperature was decreased and held at 23 °C. Further review of simulation results showed that keeping a lower setpoint during these hours required operating two chillers. In this building, the chilled water plant contains two identical chillers controlled to uniform loading when both are operating. Since FR is provided through chilled water setpoint modulation in the large office, nearly double the FR response can be achieved when both chillers are active and

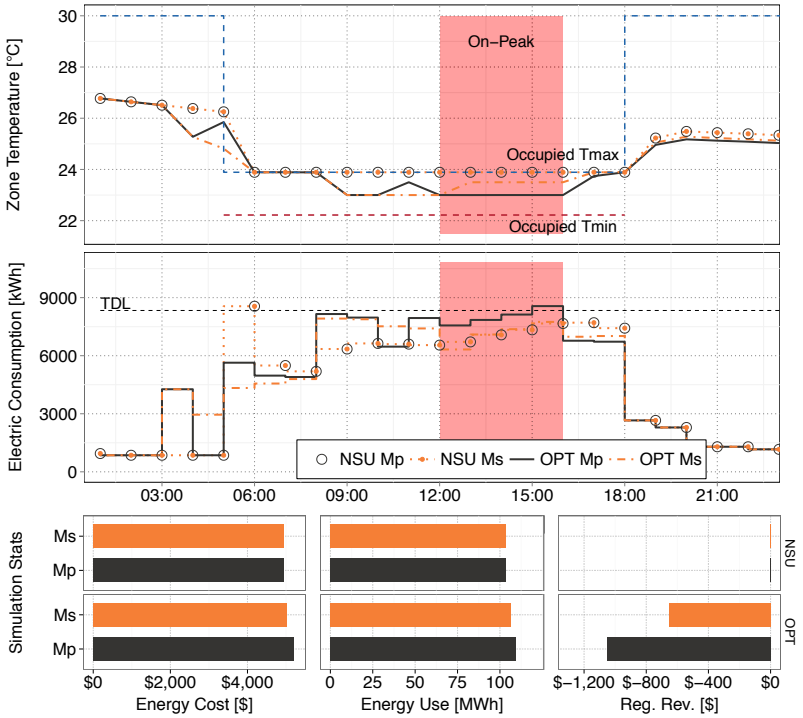

Figure 9.5: Case 3a medium office (x20) optimization results including FR.

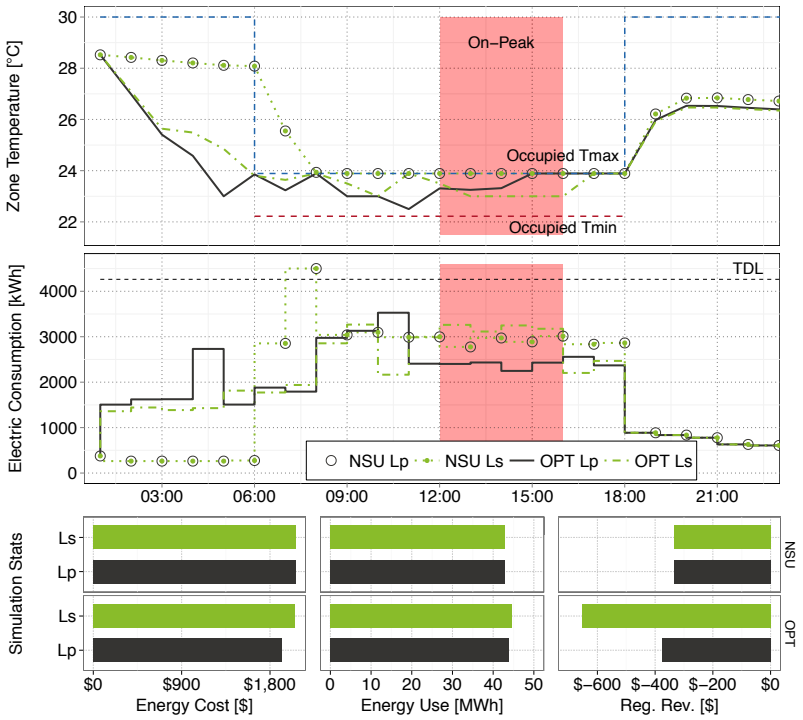

Figure 9.6: Case 3a large office (x1) optimization results including FR.

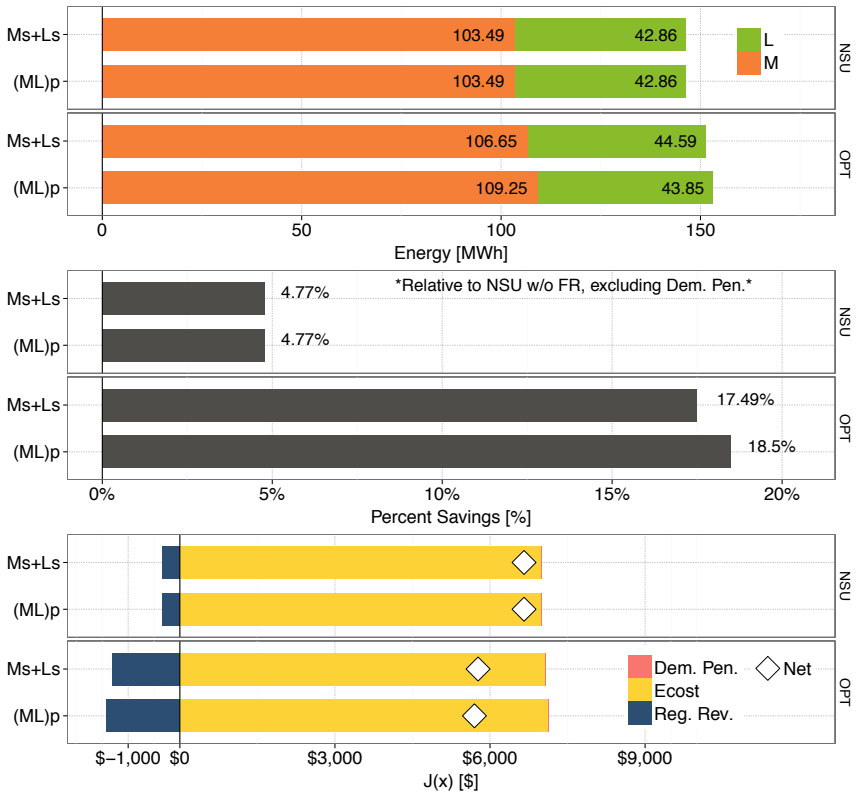

Figure 9.7: Case 3a aggregated individual and portfolio optimization results with FR.

Table 9.3: Case 3a aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	333.50	1301.90	333.50	1430.04
Energy Cost [$]	6992	7071	6992	7128
Energy Use [kWh]	146 352	151 246	146 352	153 102
Demand Penalty [$]	0	0	0	0
$J(\vec{x})$ [$]	6658	5769	6658	5698
% Difference	-4.77%	-17.49%	-4.77%	-18.50%

modulating in unison. High regulation prices were observed for the hours of 1:00 PM to 4:00 PM, thus increasing the FR response during this period resulted in a significant increase in regulation revenue. The OPT Ls case required nearly the same energy expense as the NSU Ls case, despite consuming 4% more energy. The OPT Lp solution shows significantly less regulation revenue and much lower baseline demand during peak hours. This suggests that it may be more favorable to provide FR using the medium office buildings.

The aggregate results for scenario "a" with FR are summarized in Figure 9.7. The NSU cases were identical in energy consumption, optimizer cost, and percent savings implying that no natural synergy existed through aggregation under NSU control. The large office NSU case was able to generate $334 in regulation revenue, which translates into 4.77% savings. The OPT (ML)p case consumed 1856 kWh (1.2%) more than the OPT Ms+Ls, costing an additional $57 (0.8%) in energy expense. However, the portfolio optimization was able to generate $128 (9.8%) more FR revenue, resulting in an overall increase in savings of one percentage point. The cost function details are numerically summarized in Table 9.3 for reference.

9.2 Case 3b: Both Buildings Peak Constrained

Scenario "b" represents the situation when both building have peak demand constraints and are unable to attain their maximum regulation strategy. This situation may arise when both buildings have already set their peaks for the month and are currently operating below. This scenario was simulated by applying target demand limits of 8880 kW and 3720 kW to the medium and large office buildings, respectively. Unlike the scenario "b" in Chapters 7 and 8, the TDL in this case was not evenly distributed between buildings. Approximately 73% was assigned to the medium office while 27% was assigned to the large office. This distribution was chosen in order to allow for the NSU cases to provide maximum FR without incurring demand penalties.

9.2.1 No Frequency Regulation

Figures 9.8 – 9.10 highlight the results of scenario "b" optimizations excluding frequency regulation. As in scenario "a," the target demand limits do not interfere with operation when FR is not included and the optimizations essentially seek to reduce energy expense. Therefore, the results were identical to scenario "a," and further discussion is omitted. Cost function details are provided for reference in Table 9.4.

9.2.2 With Frequency Regulation

Scenario "b" was repeated with the inclusion of frequency regulation in the optimization. Figure 9.11 shows the optimization results for the medium office building. The OPT Ms strategy began conditioning two hours before the NSU case, and kept the zone temperature near the middle of the temperature bounds from 9:00 AM to 4:00 PM to create opportunities for frequency regulation. At first glance, the OPT Ms strategy appears nearly identical to that of scenario "a," however a closer look reveals that the optimizer kept the zone temperature at 23 °C during hours ending 1:00 PM and 2:00 PM (i.e. rather than raising to 23.5 °C). Maintaing a lower temperature during these hours allowed for \$76 (11.7%) more regulation revenue while not exceeding the demand threshold. Compared to the NSU Ms case, the OPT Ms generated \$726 in regulation revenue while only consuming 4% more energy and requiring \$139 (2.8%) more energy expense. In general, this setpoint strategy appears closer to the OPT Mp solution.

The large office results for scenario "b" including regulation are provided in Figure 9.12. Compared with scenario "a" the OPT Ls case resulted in a later start-up and warmer zone temperatures. The lower large office TDL in scenario "b" constrained the OPT Ls case to approximately half of the on-peak FR capability observed in scenario "a," resulting in \$160 (25%) less regulation revenue. The warmer average zone temperatures did, however, allow the OPT Ls case to consume 1703 kWh less (4%) and spend \$60 (3%) less on energy compared to OPT Ls in scenario "a." The OPT Lp strategy performed more precooling and kept temperatures cooler on average. Less regu-

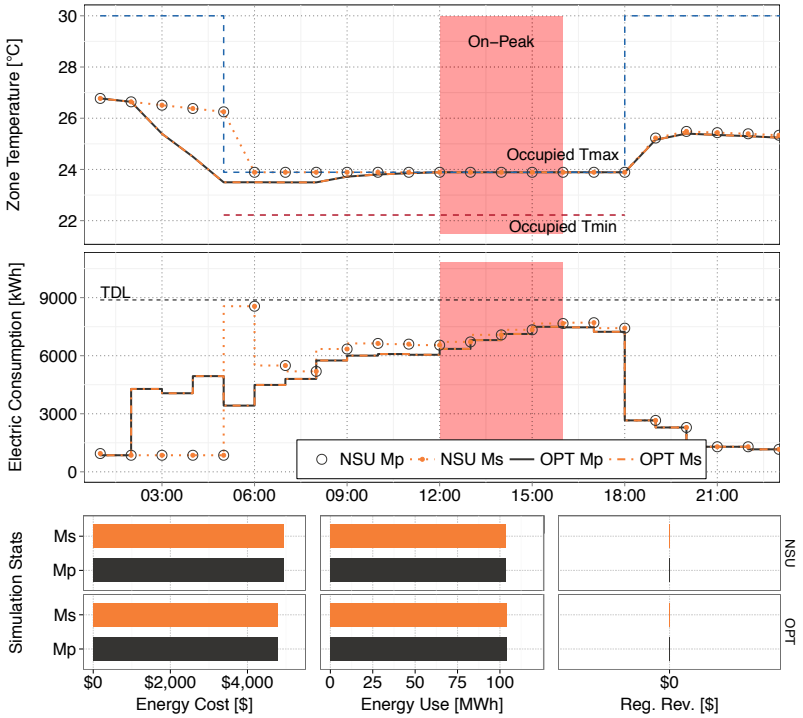

Figure 9.8: Case 3b medium office (x20) optimization results.

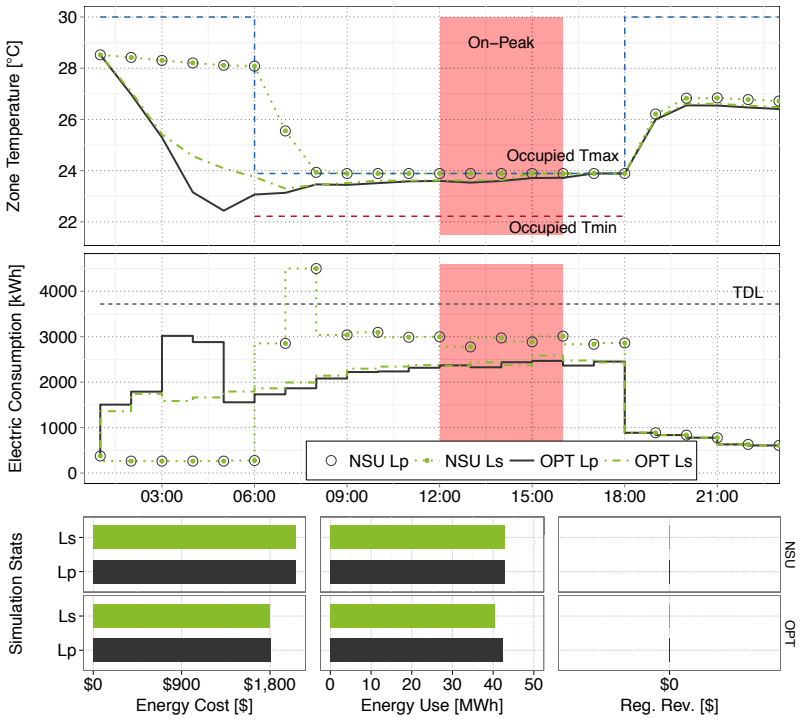

Figure 9.9: Case 3b large office (x1) optimization results.

Figure 9.10: Case 3b aggregated individual and portfolio optimization results.

Table 9.4: Case 3b aggregated individual and portfolio optimization results.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	0	0	0	0
Energy Cost [$]	6992	6593	6992	6605
Energy Use [kWh]	146 352	144 532	146 352	146 336
Demand Penalty [$]	0	0	0	0
$J(\vec{x})$ [$]	6992	6593	6992	6605
% Difference	0%	-5.70%	0%	-5.53%

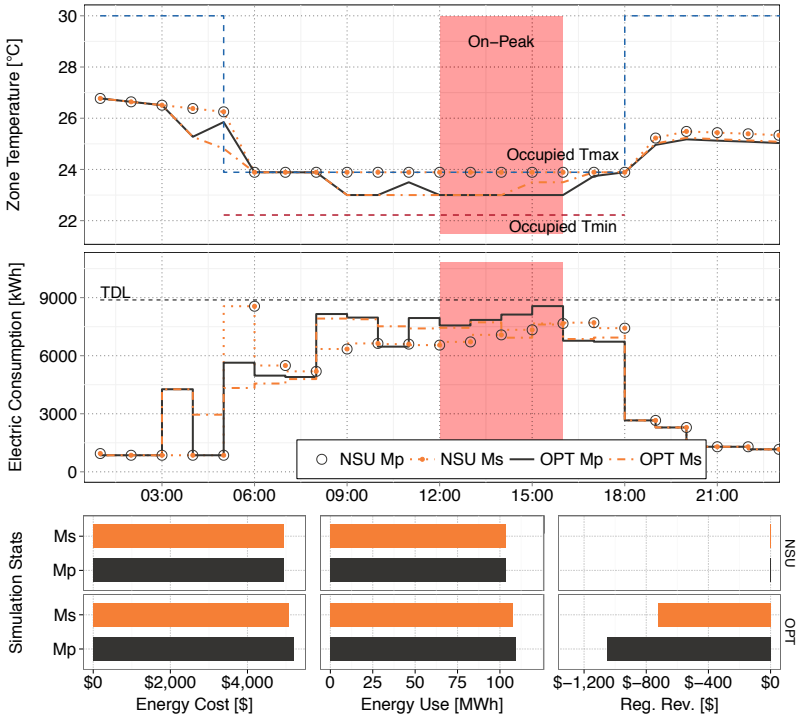

Figure 9.11: Case 3b medium office (x20) optimization results including FR.

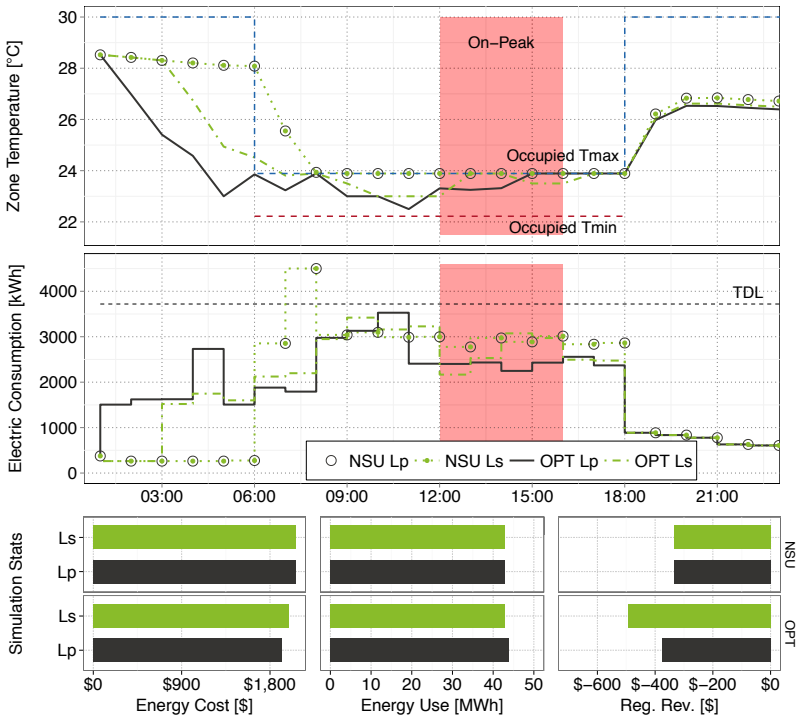

Figure 9.12: Case 3b large office (x1) optimization results including FR.

lation revenue was generated by the OPT Lp case, however energy expense savings were observed over the OPT Ls case.

The aggregate results for scenario "b" including frequency regulation are shown in Figure 9.13. Comparing the aggregated individual and portfolio NSU cases shows that energy consumption and optimizer cost are identical, implying that no diversity exists, in this scenario, when aggregated under NSU control (i.e. individual peaks occur coincidently). Simply, adding frequency regulation to the NSU operation allowed for 4.77% total cost savings. Comparing the energy use of the optimal cases shows that the OPT (ML)p strategy consumed 2585 kWh (1.7%) more than the OPT Ms+Ls case. The additional energy consumption resulted in only \$71 (1%) more energy expense, but enabled an additional \$212 (17%) in regulation revenue. Overall, the portfolio optimization achieved two percentage points more savings than the aggregate individual optimizations. A numerical summary of the cost function values is provided in Table 9.5.

9.3 Case 3c: Large Office Peak Constrained

Scenario "c" represents the situation where the large office is peak constrained while the medium office is able to attain its maximum frequency regulation operation. This situation may arise due to variations in operational patterns and control strategies that cause buildings to peak on different days. In this scenario, the large office may be experiencing a peak-limiting day, while the medium office has already set a peak higher than the current operating day. This scenario was simulated by applying target demand limits of 10 340 kW and 2260 kW to the medium and large office buildings, respectively.

9.3.1 No Frequency Regulation

Scenario "c" results for the medium office building are shown in Figure 9.14. Since the target demand limit was much higher than the NSU demand and FR was not included in this case, the optimal strategies are only incentivized by hourly energy prices. These results are identical to those observed in scenario "a" and "b" when regulation was not included.

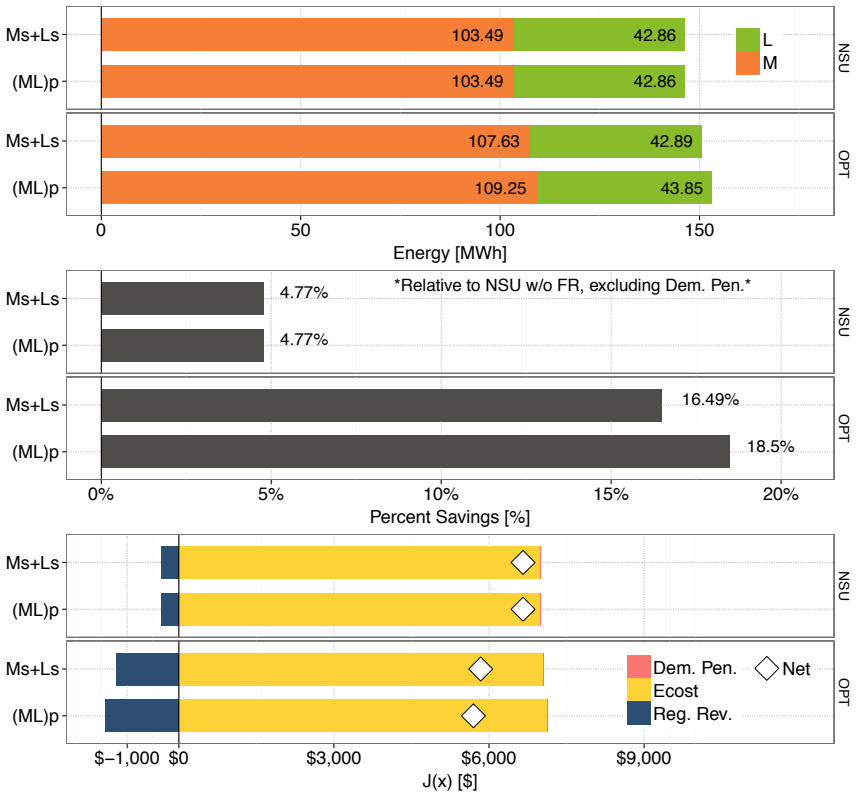

Figure 9.13: Case 3b aggregated individual and portfolio optimization results with FR.

Table 9.5: Case 3b aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	333.50	1218.33	333.50	1430.04
Energy Cost [$]	6992	7057	6992	7128
Energy Use [kWh]	146 352	150 516	146 352	153 102
Demand Penalty [$]	0	0	0	0
$J(\vec{x})$ [$]	6658	5839	6658	5698
% Difference	-4.77%	-16.49%	-4.77%	-18.50%

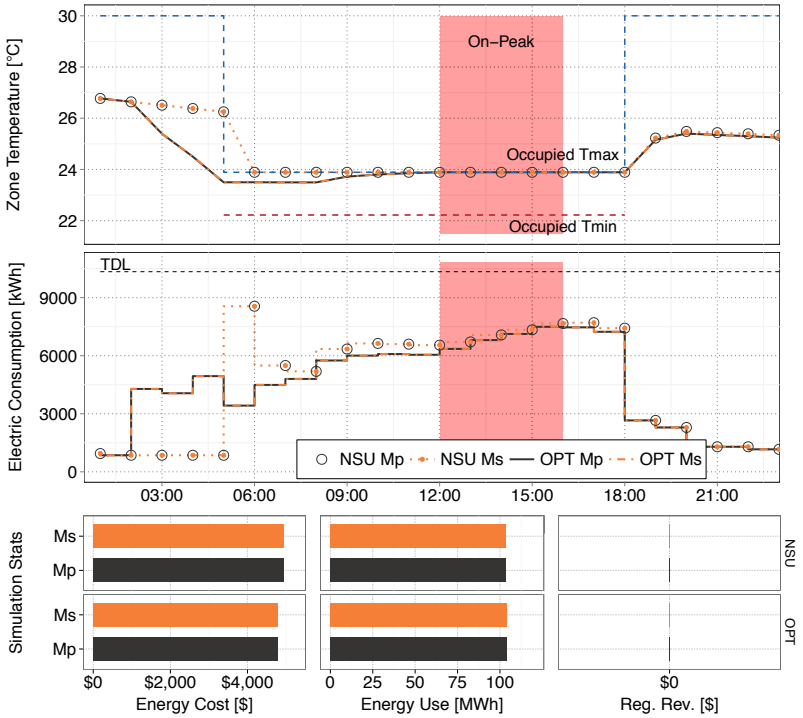

Figure 9.14: Case 3c medium office (x20) optimization results.

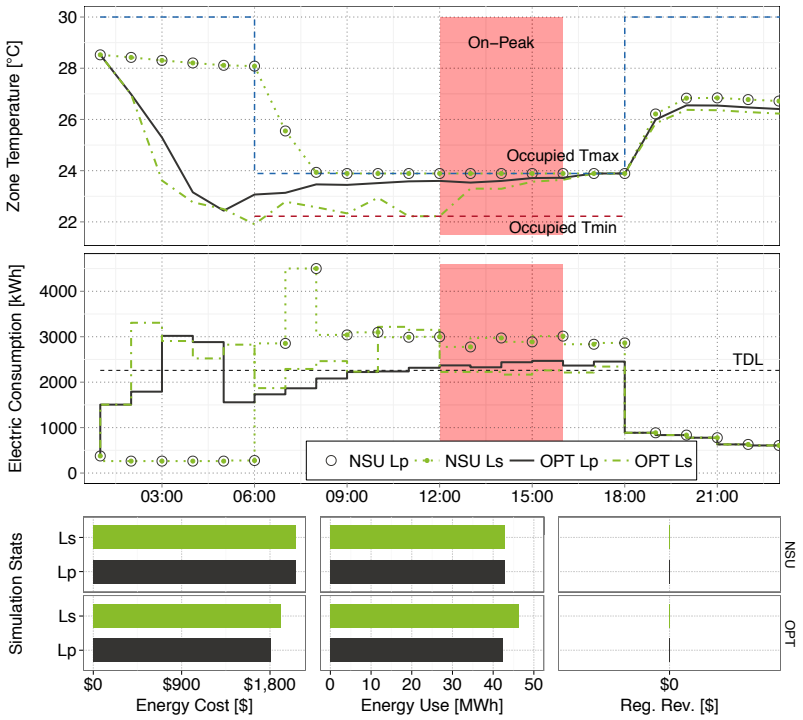

Figure 9.15: Case 3c large office (x1) optimization results.

The large office building optimization results for scenario "c," excluding regulation are shown in Figure 9.15. The large office TDL was below the NSU demand profile requiring the OPT Ls case to reduce peak demand in order to avoid costly demand penalties. The optimizer chose to begin conditioning at 1:00 AM and maintain zone temperatures near the lower bound until the beginning of the on-peak period. The zone temperatures were then raised to the upper bound to achieve the necessary demand reductions. Avoiding expensive demand penalties required the OPT Ls case to consume 3572 kWh (8.3%) more energy than the NSU Ls case, costing an additional $150 (7.3%) in energy expense. Comparatively, the OPT Lp strategy performed less precooling, although demand was significantly reduced from the NSU peak.

The aggregate results for scenario "c" are shown in Figure 9.16. The OPT (ML)p strategy achieved the portfolio demand target while consuming 4093 kWh (2.7%) less energy than the OPT Ms+Ls case and spending $103 (1.5%) less in energy expense. The energy savings is attributable to the fact that the portfolio optimization can observe that the medium office is nowhere near its demand limit and determine that it is unnecessary for the large office to reduce its peak. The overall savings in this example were computed excluding the demand penalty since this was the only scenario that incurred a demand penalty and it is eventually desired to compare savings across all scenarios. Cost function values are summarized in Table 9.6.

9.3.2 With Frequency Regulation

Scenario "c" optimizations were also repeated including frequency regulation. Figure 9.17 highlights the optimization results for the medium office building. As previously noted, this configuration allowed the medium office building to attain its maximum regulation solution. This is observed in the zone temperature being held at the middle of the temperature bounds from 9:00 AM to 4:00 PM to create FR capability. The OPT Ms case generated $1101 in regulation revenue, while spending $275 (5.6%) more than the NSU Ms case on energy. The increased energy expense was attributable to 5% higher energy consumption required by the OPT Ms case to maintain the lower zone temperatures. The optimal portfolio solution for the medium office building shows a nearly

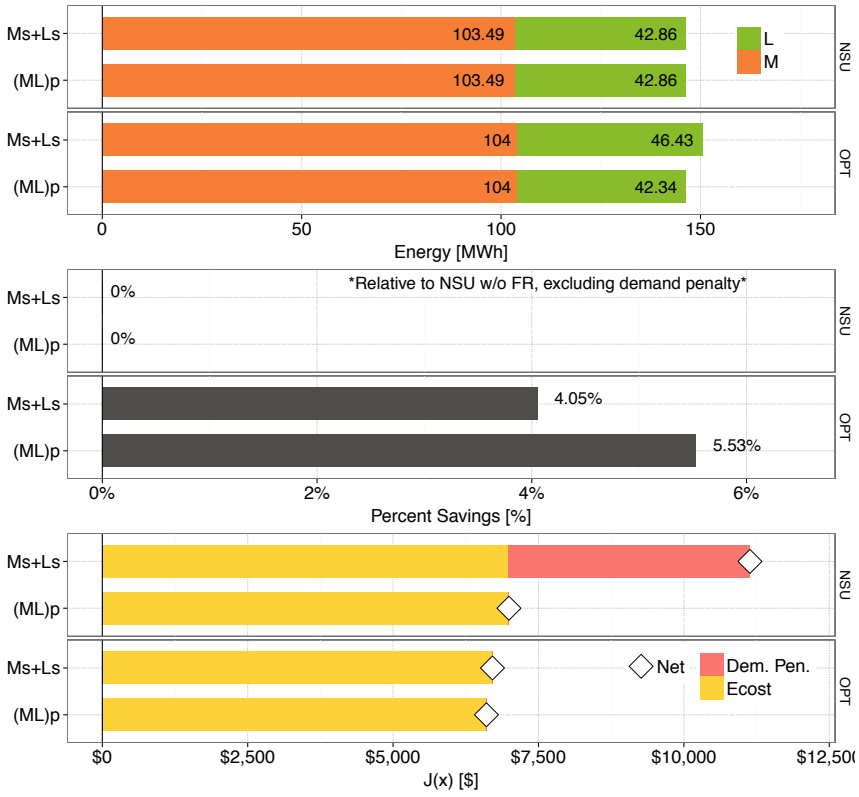

Figure 9.16: Case 3c aggregated individual and portfolio optimization results.

Table 9.6: Case 3c aggregated individual and portfolio optimization results.

	Ms+Ls			(ML)p	
	NSU	OPT		NSU	OPT
Reg. Revenue [$]	0	0	0	0	
Energy Cost [$]	6992	6708		6992	6605
Energy Use [kWh]	146 352	150 429		146 352	146 336
Demand Penalty [$]	4143.92	0		0	0
$J(\vec{x})$ [$]	11 136	6708		6992	6605
% Difference*	0%	-4.05%		0%	-5.53%

*Computed excluding demand penalty

identical zone temperature strategy and demand profile. The OPT Mp case generated $48 (4.3%) less regulation revenue than the OPT Ms case, while energy use and energy expense remained nearly the same (i.e. 0.5% and 0.1% difference in energy use and energy expense, respectively).

The large office results for scenario "c" with frequency regulation are provided in Figure 9.18. Comparing NSU results shows that the TDL in the NSU Ls case resulted in less FR capability than the portfolio case. As noted in the non-FR case, the low TDL also caused the OPT Ls case to perform significant precooling in effort of avoiding significant demand penalties. This required consuming 5006 kWh (11.7%) more energy than the NSU Ls case, however more energy use was shifted to lower priced hours and $81 (4%) of energy expense savings were observed. The OPT Ls case was able to generate $285 of regulation revenue, 38% more than the NSU Ls case. In general the portfolio solution resulted in less precooling, more FR, and less energy expense compared to the OPT Ls strategy.

The aggregate results for scenario "c" are provided in Figure 9.19. Although the NSU cases consumed an equal quantity of energy, the NSU Ms+Ls case was able to perform less regulation revenue and incurs a demand penalty due to the low target demand. Therefore, aggregating under NSU control in this example resulted in saving an additional 1.8 percentage points over the NSU Ms+LS case. Comparing the OPT results shows that the portfolio optimization resulted in 3435 kWh (2.2%) less energy consumption, $56 (0.8%) less energy expense, and $44 (3.2%) more regulation revenue than the aggregated individual optimizations. Overall this translates into 1.45 percentage points of additional savings. A numerical summary of the simulation cost functions is provided in Table 9.7.

9.4 Case 3 Summary

Comparing the results across scenarios "a", "b", and "c" for Case 3, highlights that the magnitude of observed synergy, again, may be somewhat dependent on the conditions each building would have experienced before aggregation. Table 9.8 summarizes the achieved percent savings by the OPT cases for the investigated scenarios. The demand penalties were excluded in the percent

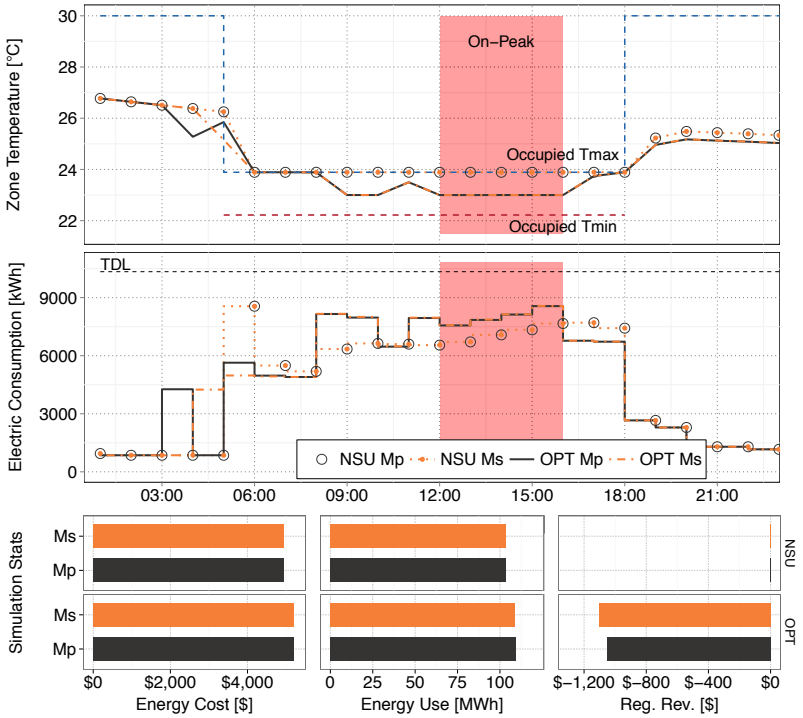

Figure 9.17: Case 3c medium office (x20) optimization results including FR.

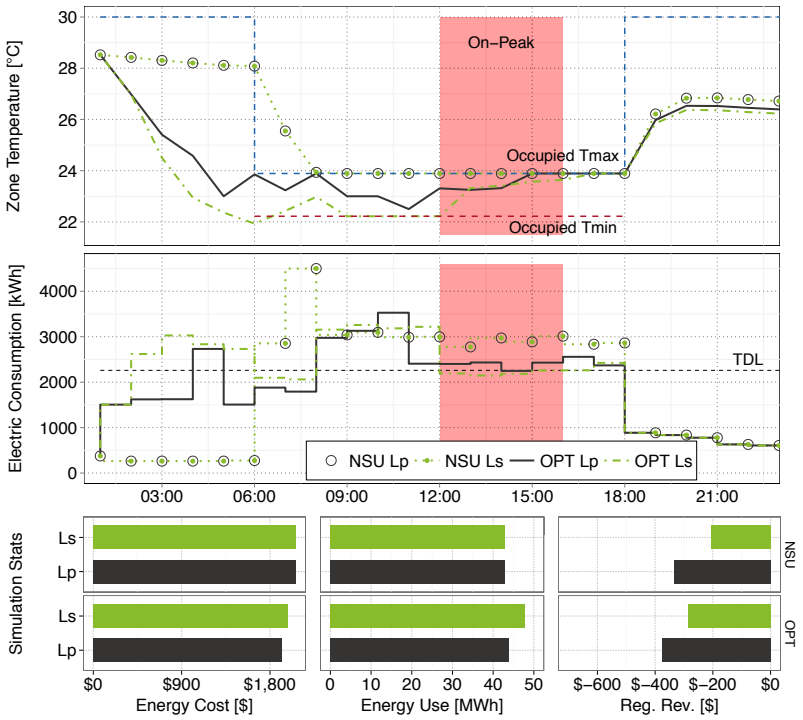

Figure 9.18: Case 3c large office (x1) optimization results including FR.

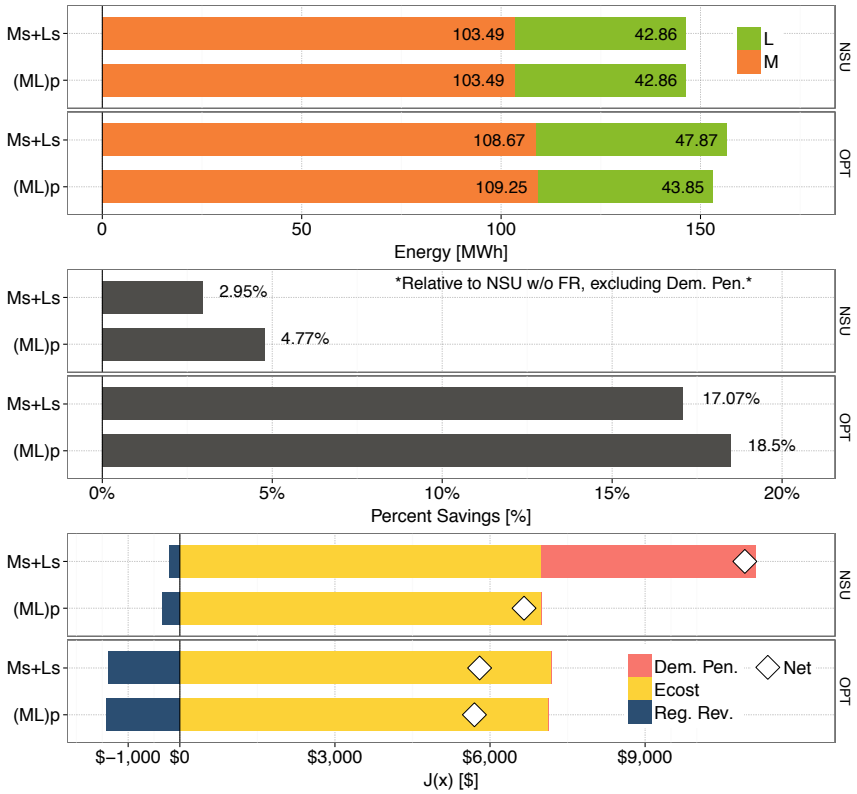

Figure 9.19: Case 3c aggregated individual and portfolio optimization results with FR.

Table 9.7: Case 3c aggregated individual and portfolio optimization results with FR.

	Ms+Ls		(ML)p	
	NSU	OPT	NSU	OPT
Reg. Revenue [$]	206.14	1385.87	333.50	1430.04
Energy Cost [$]	6992	7184	6992	7128
Energy Use [kWh]	146 352	156 537	146 352	153 102
Demand Penalty [$]	4143.92	0	0	0
$J(\vec{x})$ [$]	10 929	5798	6658	5698
% Difference*	-2.95%	-17.07%	-4.77%	-18.50%

*Computed excluding demand penalty

savings calculation of scenario "c" so that all NSU Ms+Ls cases were identical and a fair comparison could be made across scenarios.

Table 9.8: Case 3 percent savings summary.

	No FR			FR		
Scenario	Ms+Ls	(ML)p	Diff.	Ms+Ls	(ML)p	Diff.
a	-5.70%	-5.53%	-0.17	-17.49%	-18.50%	1.01
b	-5.70%	-5.53%	-0.17	-16.49%	-18.50%	2.01
c	-4.05%	-5.53%	1.48	-17.07%	-18.50%	1.43

As expected, the results without frequency regulation were somewhat trivial since the high TDL had little influence on operations without FR. One exception was observed in scenario "c" where the portfolio optimization resulted in nearly 1.5 percentage points more savings compared to the aggregate individual optimizations. However, this synergy was similar to that observed in Case 1 and Case 2 where the portfolio perspective can avoid unnecessary precooling due to having knowledge that the buildings do not exceed the demand limit when aggregated.

Considering the results with frequency regulation, the portfolio optimization determined that for this specific portfolio construction (i.e. 20 medium office, and one large office) and simulation case the optimal operation was achieved when the medium office operates near its maximum regulation solution and the majority of demand reductions are performed by the large office building. The conditions of scenario "c" were most similar to this, however 1.43 percentage points of additional savings were still observed due to the fact that diversity among the portfolio allowed for generation of more regulation revenue.

Slightly less synergy was observed for scenarios where one building provides max FR while the other provides demand relief (i.e. "a" and "c"). This suggests that splitting regulation and peak reducing duties was preferable over constraining both buildings. This result seems plausible since it was previously observed that either building could make significant demand reductions by

responding to hourly energy prices. If one building pursues opportunities for energy arbitrage and efficiency gains, the resulting demand reductions would empower another building to seek regulation maximizing strategies.

In this case, the medium office buildings provided frequency regulation through zone temperature modulation which excites a power response from the whole HVAC system. The large office provided FR via chilled water setpoint modulation, exciting a power response from the chillers only. Comparing the maximum regulation scenarios between the large and medium office buildings shows that the 20 medium office buildings have approximately double the FR capability of the single large office building. Furthermore, previous optimizations showed that the large office building was able to make larger voluntary demand reductions than the medium office in pursuit of energy arbitrage. Therefore, it also seems plausible that greater benefit would be achieved by dispatching the medium office towards maximizing regulation while the large office provides peak demand relief.

Although interesting synergies were observed, the overall percent savings increases were somewhat small at 1 to 2 percentage points. Providing more FR generally required consuming more energy, thus there were no beneficial end-use electricity savings to be weighed along with the cost savings. In this specific case, aggregating the individual optimizations performed quite well and the additional complexity of a full portfolio optimization may be undesirable for the magnitude of benefit observed.

Chapter 10

Conclusions and Discussion

10.1 Summary

In Chapter 1, this research was motivated by discussing the growing and changing global demand for energy. More energy will be demanded in the future as nations progress towards higher levels of development, and a greater proportion of energy demand will be shifted towards modern energy sources and services (e.g. electricity). Furthermore, buildings will play an increasing role in global energy demands as developing nations further advance their stock of modern facilities. Satisfying growing energy demands in a sustainable manner requires advancing the efficiency of current technology, as well as increasing the adoption and utilization of renewable energy resources. Improved integration of commercial building operations with the electric grid provides an opportunity to enable variable renewable generation by creating demand flexibility. Additional benefits may be possible through controlling groups of buildings from a higher-level portfolio perspective.

A review of literature surrounding relevant topics to this research was provided in Chapter 2.

To enable exploration of research questions regarding optimization of building portfolios participating in multiple grid markets, a reduced-order modeling and simulation environment was developed and discussed in Chapter 4. Building and HVAC models for retail, medium office, and large office buildings were developed and validated with respect to detailed EnergyPlus models. The reduced-order models showed speed improvements of approximately one to two orders of magnitude over optimizations with EnergyPlus, while maintaining adequate accuracy for the desired simulation

studies.

As a first step towards better integration of commercial buildings with the electric grid, methods to estimate the capability of a commercial building to provide frequency regulation ancillary services were developed in Chapter 5. Overall, the methodology presented seems promising for estimating commercial building regulating potential since detailed interactions and constraints are captured through the use of whole-building energy models. The MPC framework, considering both cost and revenue, appropriately determined opportunities when buildings could economically provide frequency regulation. The optimizer was able increase the potential regulation capability by moving buildings away from setpoint and system capacity limits to a more flexible state. Additionally, the modeling process generated further insight into providing FR through zone and CHW setpoints. As expected, modulating the zone setpoint can excite an organized response from the whole HVAC system. When injecting the FR signal at lower system levels, additional compensating effects may occur that limit the potential magnitude and duration of a FR response. Overall, including frequency regulation in the optimization increased the total savings by one to 13 percentage points, depending on the case, creating potentially significant benefits for both building owners and the electric grid.

In Chapter 6, the multi-market optimization was extended to accommodate building portfolios. Preliminary portfolio optimizations were performed to gain confidence in the optimization algorithm as well as inform the design of further portfolio optimization studies.

Chapter 7 presented the optimization results for a portfolio of 120 retail buildings and one large office. Depending on the conditions of the individual optimizations, 0.5 to 7 percentage points of additional total cost savings were observed when buildings were optimized as a portfolio. Portfolio optimizations were able to satisfy demand limits while consuming 1.6% to 17.5% less energy than the aggregated individual optimizations, due to the fact that the retail buildings were quite inefficient at shifting load.

Portfolio optimization results for 20 medium office and one large office building were presented in Chapter 8. Slight synergistic effect of one percentage point was observed through portfolio

optimization, due to the fact that voluntary demand reductions made by each building in response to real time pricing incentives could be allocated to the total desired portfolio demand reduction. The portfolio optimization also resulted in 3% less electricity consumption in one case as well.

A third portfolio optimization case study was presented in Chapter 9, again consisting of 20 medium office buildings and one large office building. One to two percentage points of additional savings were observed for the portfolio optimizations, depending on the conditions of the individual optimizations. Synergy was observed through the large office building taking on peak reducing responsibilities to enable the medium office building to pursue greater amounts of frequency regulation.

10.2 Conclusions

Overall, this work has demonstrated the potential benefit of co-optimizing building participation in multiple grid markets. Furthermore, it was observed that synergistic effect can be achieved through optimization of building portfolios, although the magnitude and nature of synergistic effect may be case dependent. From the simulation studies performed it can be concluded that synergistic effect may result from a variety of pathways.

First, synergy may arise if a building is better at performing a specific task. In Portfolio Case 1, the thermal mass and envelope properties of the large office building made it inherently better than the retail buildings at passive thermal storage. Energy savings were observed when the large office building provided additional load reductions to account for the retail buildings target demand limits. This may suggest that the large office building should provide additional load reductions whenever possible, and that the retail buildings should only shift load when the large office building does not have the extra demand reducing capacity (i.e. peak days). This may also suggest that controlling one thermally massive building can account for the unoptimized operation of multiple inefficient buildings. In Portfolio Case 3, the large office could make larger voluntary demand reductions in response to energy pricing incentives, while the medium office building had greater capacity for providing FR. The optimization therefore determined that it was better to

assign load reducing duties to the large office building and allow the medium office to pursue the opportunity to maximize FR participation.

Second, portfolio diversity may allow attainment of objectives at lower cost and energy. In other words, portfolios may not have to work as hard to achieve the same aggregate objectives. In Portfolio Case 2, the portfolio perspective allowed both buildings to utilize voluntary demand reductions made in response to real time pricing incentives towards peak reducing objectives. Individually, buildings were required to achieve greater load reductions since the voluntary load reductions of the other building were unknown.

Third, portfolio diversity may allow for greater FR. In seven of the nine scenarios investigated, the portfolio optimizations resulted in greater amounts of FR. Increased frequency regulation can result from the scenario where one building has excess FR capability, but is constrained by its demand limit; while the other building has space below its demand limit to provide more FR, but does not have the capability. Additionally, the multi-market optimizations showed that a buildings FR capability is strongly influenced by the thermal mass strategy since the trajectory informs where and when equipment and operational constraints may occur. It was observed that when a building is required to perform extreme load reductions it is often operating near setpoint or equipment limitations and, in general, has less flexibility to provide FR.

In addition to conclusions regarding how synergy may arise within building portfolios, the simulation studies performed in this research may be used to draw conclusions about the nature of synergistic effect in building portfolios as well.

First, synergistic effect may be dependent on the individual optimization conditions. In all portfolio optimizations the magnitude of observed synergy was dependent on the target demand limit of the individual building optimizations. If buildings were relatively similar as in Portfolio Case 2, and tended to peak on the same days (i.e. have similar peak shaving requirements), little opportunity for synergy was observed. However, if those same buildings peak on different days, synergistic effect may be possible.

Second, synergistic effect may be dependent on portfolio construction. Varying amounts of

synergy from a variety of sources were observed in the portfolio optimizations. The least amount of synergy was observed in Portfolio Case 2 where the buildings were relatively similar, and the greatest amount of synergy was observed in Portfolio Case 1 where buildings had significantly different characteristics.

Third, synergistic effect may be dependent on grid market design. When peak demand was excluded from the optimizations there was no opportunity for synergy. For synergy to exist among building portfolios a joint objective must be in place, which may ultimately be defined by rate structures and the opportunities available within grid markets.

10.3 Discussion, Limitations, and Future Work

Overall, this research supports the hypothesis that buildings can work cooperatively towards joint objectives when given the appropriate motivation. The common motivation used in this work was the management of a communal peak demand. The portfolio target demand is essentially considered a monthly ceiling that the portfolio tries to avoid exceeding. Synergistic effect was observed on non-peak setting days when buildings have the flexibility for a variety of operational strategies below this demand ceiling. However, rather than adhering to the monthly demand ceiling in a more efficient way, the buildings could have contributed larger load reductions to the grid if desired and incentivized. It may actually be more beneficial to the grid if buildings would contribute maximum demand reductions every day. This brings into question the role of monthly peak demand charges in a modern electric grid where it may be more beneficial to focus on real-time signals that communicate the current operational desires of the grid.

The dependence of synergy on a peak demand connection between buildings in this work can be considered a limitation since no synergy would have been observed in its absence. The majority of synergy observed under the demand limiting setting could generally be described as one building having a desire to avoid a demand limit and another building being capable of assisting in avoidance. In all cases the buildings could have avoided the demand limit through their own operation, although it may have been beneficial to receive assistance from another building. It

seems possible that another category of incentives to cooperate may exist when buildings are given tasks that they cannot perform alone. One example of this could be if the grid desires a staged or ramped building demand profile to accommodate generation ramping. In general the question of what joint objectives exist for building portfolios in a future smart grid seems paramount, and remains largely unanswered by this research. Future work could re-consider portfolio optimization in conjunction with novel energy market designs that create new opportunities for cooperative load control.

Aside from further exploration of opportunities that motivate a cooperative response from buildings, a myriad of related questions exists surrounding the topic of building portfolio optimization. First, this work considered the optimization within the context of a single climate and pricing scenario. Further analysis should be performed to better understand how the opportunities for synergy materialize under more complex peak rate structures and their sensitivity to pricing signals.

Second, portfolios were constructed somewhat arbitrarily in this work to investigate the mere existence of synergistic effect. However, the aggregation process could be considered an optimization problem in and of itself to maximize the portfolio synergy. This work has only scratched the surface of exploring how diversity among building characteristics may translate into synergistic effect within a portfolio, and future work could seek to rank the importance of various building parameters in leading to synergy (see Appendix A). Building portfolio optimizations could also be framed to incorporate risk management as is common with financial portfolios.

Third, the analysis should be extended to include full month-long simulation studies to evaluate how often scenarios (such as a, b and c) arise that are favorable for portfolio synergy. No attempt was made in this work to quantify the long-term opportunities for savings through synergistic control.

Fourth, the optimization may benefit from a multi-objective formulation. In the current framework, output variables of interest (i.e. energy cost, peak demand penalty, regulation revenue, etc.) are all combined within the cost function to arrive at a single value for the control strategy.

Past experience has shown that the results are often sensitive to the assignment of penalty terms and the magnitude of potentially fictitious dollar amounts in the cost function. Taking a multi-objective optimization approach would allow greater exploration of the trade-offs present among cost function dimensions and perhaps allow for more informed decision making.

Perhaps one of the most significant limitations of this work is that no opportunity was available for field testing of new methods. The portfolio optimization is essentially an extension of the previously tested single building MPC controller and it may seem reasonable that the portfolio extension should perform similarly in a field setting. However, the frequency regulation estimation methods remain untested outside of simulation and future work should seek to validate the methods with respect to field data. Section 5.5 provides some discussion of the implications of modeling on the frequency regulation estimation, however, in general it is not understood how well hourly model perturbation results capture the flexibility available in an actual commercial building. Frequency regulation markets often require hourly commitment of frequency regulation capability, thus, an attempt at an hourly analysis may seem reasonable as a first approach for planning thermal mass strategies in consideration of frequency regulation opportunities. However, the signal is ultimately dispatched at much shorter 2 to 4 second timescales and it may be necessary to develop models or methods that characterize the short-term dynamic response of a building in order to provide accurate estimates. Furthermore, the regulation signal is ultimately random and the characteristics of the signal that materialize in real-time may prove to have the greatest influence on the magnitude of the response that can be provided by a building. It is difficult to know whether or not (or how much) the estimation methodology could benefit from a full dynamic description of the building if the regulation signal cannot be predicted in advance. Future work could investigate the simulation resolution and model dynamics required to accurately predict building FR flexibility. Opportunities may also exist for stochastic analyses due to the random nature of the FR signal.

Both the centralized portfolio optimization and model perturbation FR estimation approaches used in this work are plagued by practicality considerations that must be addressed before wide-spread implementation is possible. Considering all decision variables for all buildings means that

the optimization problem becomes intractable as the size of the portfolio grows. The potential for scalability was included in this work through the inclusion of a building multiplier. However, the extendability of optimal control strategies developed for one building to other similar buildings has yet to be explored, and may ultimately prove unreliable given the diverse nature of buildings. The complexity can be reduced by taking a courser discretization of the control vectors, however, even so the approach is still likely limited to relatively small portfolios. Although simplifications were made with respect to the building thermal model, it was assumed in this work that a relatively high degree of complexity was required in the HVAC models to characterize the building demand profile for large commercial buildings and develop adequate control strategies. Future work should explore model complexity questions with respect to portfolio optimization and seek to determine whether simpler models or meta-modeling strategies are appropriate. It may be determined that when considering the response of thousands of buildings, model mismatch and inadequacies may be less critical than when evaluating the optimal response of a single building. The FR estimation would benefit from model simplification as well since the current approach runs potentially hundreds of simulations for each MPC function evaluation. More elegant solutions to FR estimation should also be sought that further reduce the simulation burden.

In general, the work performed in this research seems to illuminate more complexities and questions, than produce generalizable results. Although much room for expansion and improvement exists, the tool developed in this work can be used in its current form to answer many of the questions previously discussed through offline simulation studies. The offline optimizations could be utilized to glean novel heuristics for portfolio construction and synergistic control for application in real scenarios (see Appendix A). Machine learning or data mining methods could be applied to extract potentially non-trivial relationships that exist among the simulation data.

As buildings become more active participants in the smart grid and building-integrated energy resources increase in adoption, the previously distinct line between supply and demand sources will become blurred. Advanced control paradigms will be critical at both building and grid levels to achieve harmonious integration and a sustainable energy future. An interesting aspect of

this research is that the general theory surrounding MPC has remained largely intact, while the controller perspective has been shifted from an individual to a communal perspective. Additional benefits were observed from this change in perspective, which may motivate the pursuit of future smart grid advancements that take a communal or holistic vantage point.

Bibliography

[1] Athanassios A Argiriou, Ioannis Bellas-Velidis, Michaël Kummert, and Philippe André. A neural network controller for hydronic heating systems of solar buildings. Neural Networks, 17:427–440, 2004.

[2] Peter R Armstrong, Steven B Leeb, and Leslie K Norford. Control with Building Mass Part I : Thermal Response Model. ASHRAE Transactions, 112(1), 2006.

[3] Peter R Armstrong, Steven B Leeb, and Leslie K Norford. Control with Building Mass Part II : Simulation. ASHRAE Transactions, 112, 2006.

[4] ASHRAE. ASHRAE Handbook–Fundamentals. Technical report, ASHRAE, 2009.

[5] J C Bansal, P K Singh, Mukesh Saraswat, Abhishek Verma, Shimpi Singh Jadon, and Ajith Abraham. Inertia Weight strategies in Particle Swarm Optimization. In 2011 Third World Congress on Nature and Biologically Inspired Computing, pages 633–640. IEEE, October 2011.

[6] Jonathan Berardino, Mohammed Muthalib, and Chika O Nwankpa. Network Constrained Economic Demand Dispatch of Controllable Building Electric Loads. In Innovative Smart Grid Technologies (IGST), 2012 IEEE PES, pages 1–6, Washington, DC, 2011.

[7] Jonathan Berardino and Chika Nwankpa. Dynamic Load Modeling of an HVAC Chiller for Demand Response Applications. 2010 First IEEE International Conference on Smart Grid Communications, pages 108–113, October 2010.

[8] Jonathan Berardino, Chika O Nwankpa, and Karen Miu. Economic Demand Dispatch of Controllable Building Electric Loads for Demand Response. In PowerTech, 2011 IEEE Trondheim, pages 1–6, 2011.

[9] Fatih Birol, Laura Cozzi, Tim Gould, Amos Bromhead, Christian Besson, Dan Dorner, Marco Baroni, Pawel Olejarnik, and Timur Gul. World Energy Outlook. Technical report, International Energy Agency, Paris, France, 2013.

[10] Alberto Borghetti, Claudia D Ambrosio, Andrea Lodi, and Silvano Martello. An MILP Approach for Short-Term Hydro Scheduling and Unit Commitment With Head-Dependent Reservoir. IEEE Transactions on Power Systems, 23(3):1115–1124, 2008.

[11] Cameron Bracken, Balaji Rajagopalan, and James Prairie. A multisite seasonal ensemble streamflow forecasting technique. Water Resources Research, 46(3):1–12, March 2010.

[12] M R Brambley, Phil Haves, S C McDonald, Paul A Torcellini, D Hansen, D R Holmberg, and K W Roth. Advanced Sensors and Controls for Building Applications : Market Assessment and Potential R & D Pathways (PNNL-15149). Technical Report April, Pacific Northwest National Laboratory, 2005.

[13] M.J. Brandemuehl, G. Shauna, and A. Inger. Hvac2 toolkit: Algorithms and subroutines for secondary hvac system energy calculations. ASHRAE, Atlanta, GA, 1993.

[14] James E Braun. Reducing Energy Costs and Peak Electrical Demand through Optimal Control of Building Thermal Storage. ASHRAE Transactions, 96(2):876–887, 1990.

[15] James E Braun and Nitin Chaturvedi. An Inverse Gray-Box Model for Transient Building Load Prediction. HVAC&R Research, 8(1):73–100, 2002.

[16] James E Braun, Donghun Kim, Miroslav Baric, Pengfei Li, Satish Narayanan, Shui Yuan, Eugene Cliff, John A Burns, and Bill Henshaw. Whole Building Control System Design and Evaluation : Simulation-Based Assessment. Technical Report February, GPIC Energy Efficient Buildings Hub, 2012.

[17] James E Braun, Kent W Montgomery, and Nitin Chaturvedi. Evaluating the Performance of Building Thermal Mass Control Strategies. HVAC&R Research, 7(4):403–428, 2001.

[18] H T Ceylan and G E Myers. Long-Time Solutions to Heat-Conduction Transients with Time-Dependent Inputs. Journal of Heat Transfer, 102(1):115–121, 1980.

[19] Nitin Chaturvedi and James E Braun. Analytical Tools for Dynamic Building Control (ASHRAE 985-RP). Technical Report October, Raw W. Herrick Laboratories - Purdue, 2000.

[20] Arthur I Cohen and S H Wan. An Algorithm for Scheduling a Large Pumped Storage Plant. IEEE Transactions on Power Apparatus and Systems, PAS-104(8):2099–2104, 1985.

[21] Antonio J Conejo, Miguel A Plazas, Rosa Espínola, Student Member, and Ana B Molina. Day-Ahead Electricity Price Forecasting Using the Wavelet Transform and ARIMA Models. Power Systems, IEEE Transactions on, 20(2):1035–1042, 2005.

[22] John Conti, Paul Holtberg, Joseph Beamon, Sam Napolitano, A. Michael Schaal, and James T Turnure. Annual Energy Outlook. Technical report, U.S. Energy Information Administration, Washington, DC, 2013.

[23] John Conti, Paul Holtberg, Joseph Beamon, Sam Napolitano, A. Michael Schaal, James T Turnure, and Lynn Westfall. International Energy Outlook. Technical report, U.S. Energy Information Administration, Washington, DC, 2013.

[24] Charles D Corbin, Gregor P Henze, and Peter May-Ostendorp. A model predictive control optimization environment for real-time commercial building application. Journal of Building Performance Simulation, 6(3):159–174, May 2013.

[25] Drury B Crawley, Jon W Hand, Michaël Kummert, and Brent T Griffith. Contrasting the Capabilities of Building Energy Performance Simulation Programs. Technical Report July, U S Department of Energy; University of Strathclyde; University of Wisconson-Madison; National Renewable Energy Laboratory, 2005.

[26] Drury B Crawley, Jon W Hand, Michaël Kummert, and Brent T Griffith. Contrasting the capabilities of building energy performance simulation programs. Building and Environment, 43(4):661–673, 2008.

[27] T Dewson, B Day, and AD Irving. Least squares parameter estimation of a reduced order thermal model of an experimental building. Building and Environment, 28(2):127–137, 1993.

[28] A Dhar, T Agami Reddy, and David E Claridge. Using Fourier Series to Model Hourly Energy Use in Commercial Buildings. Technical Report November, Texas A&M - Energy Systems Laboratory, College Station, 1993.

[29] Rubens A Dias, Cristiano R Mattos, and José A P Balestieri. The limits of human development and the use of energy and natural resources. Energy Policy, 34(9):1026–1031, June 2006.

[30] Robert H Dodier and Gregor P Henze. Statistical Analysis of Neural Networks as Applied to Building Energy Prediction. Journal of Solar Energy Engineering, 126(1):592, 2004.

[31] EnergyPlus Development Team. EnergyPlus-Engineering Reference. Technical report, US Department of Energy, 2010.

[32] Anthony Florita. Development of a Simulation and Optimization Environment for the Analysis of Building Thermal Mass Control. Technical report, University of Nebraska, 2007.

[33] Gilles Fraisse, Christelle Viardot, Olivier Lafabrie, and Gilbert Achard. Development of a simplified and accurate building model based on electrical analogy. Energy and Buildings, 34:1017–1031, 2002.

[34] Javier García-González, Ernesto Parrilla, and Alicia Mateo. Risk-averse profit-based optimal scheduling of a hydro-chain in the day-ahead electricity market. European Journal of Operational Research, 181(3):1354–1369, September 2007.

[35] GE Energy. Western Wind and Solar Integration Study. Technical Report May, National Renewable Energy Laboratory, Golden, CO, 2010.

[36] Pavlos S Georgilakis. Technical challenges associated with the integration of wind power into power systems. Renewable and Sustainable Energy Reviews, 12(3):852–863, April 2008.

[37] Lorna A Greening, David L Greene, and Carmen Difiglio. Energy efficiency and consumption the rebound effect a survey. Energy Policy, 28(6-7):389–401, June 2000.

[38] Erik M Greensfelder, Gregor P Henze, and Clemens Felsmann. An investigation of optimal control of passive building thermal storage with real time pricing. Journal of Building Performance Simulation Performance Simulation, 4(2), 2011.

[39] John M Halley and William E Kunin. Extinction risk and the 1/f family of noise models. Theoretical Population Biology, 56(3):215 – 230, 1999.

[40] Sekyung Han, Soohee Han, and Kaoru Sezaki. Development of an Optimal Vehicle-to-Grid Aggregator for Frequency Regulation. IEEE Transactions on Smart Grid, 1(1):65–72, 2010.

[41] He Hao, Anupama Kowli, Yashen Lin, Prabir Barooah, and Sean Meyn. Ancillary Service for the Grid via Control of Commercial Building HVAC Systems. In American Control Conference, pages 1–6, Washington, DC, 2013.

[42] Grayson Heffner, Charles Goldman, Brendan Kirby, and Michael Kintner-Meyer. Loads Providing Ancillary Services: Review of International Experience. Technical Report May, Lawrence Berkeley National Laboratory, 2007.

[43] Gregor P Henze. Energy and Cost Minimal Control of Active and Passive Building Thermal Storage Inventory. Journal of Solar Energy Engineering, 127(3):343, 2005.

[44] Gregor P Henze, Clemens Felsmann, Anthony R Florita, Michael J Brandemuehl, Hwakong Cheng, and Clarence E Waters. Optimization of Building Thermal Mass Control in the Presence of Energy and Demand Charges (1313-RP Technical Paper). ASHRAE Transactions, 114(2):75–84, 2008.

[45] Gregor P Henze, Clemens Felsmann, and Gottfried Knabe. Evaluation of optimal control for active and passive building thermal storage. International Journal of Thermal Sciences, 43(2):173–183, February 2004.

[46] Gregor P Henze, Anthony R Florita, Michael J Brandemuehl, Clemens Felsmann, and Hwakong Cheng. Advances in Near-Optimal Control of Passive Building Thermal Storage. Journal of Solar Energy Engineering, 132(2):021009, 2010.

[47] Gregor P Henze, Doreen E Kalz, Simeng Liu, and Clemens Felsmann. Experimental Analysis of Model-Based Predictive Optimal Control for Active and Passive Building Thermal Storage Inventory. HVAC&R Research, 11(2):189–213, 2005.

[48] Gregor P Henze, Thoi H Le, Anthony R Florita, and Clemens Felsmann. Sensitivity Analysis of Optimal Building Thermal Mass Control. Journal of Solar Energy Engineering, 129(4):473, 2007.

[49] IEA. Energy Efficiency Market Report: Executive Summary. Technical report, International Energy Agency, January 2014.

[50] ISO. International Standard ISO 13790. Technical report, International Organization for Standardization, 2008.

[51] S Karatasou, M Santamouris, and V Geros. Modeling and predicting building's energy use with artificial neural networks: Methods and results. Energy and Buildings, 38(8):949–958, August 2006.

[52] David Kathan. Policy and Technical Issues Associated with ISO Demand Response Programs. Technical Report July, The National Association of Regulatory Utility Commissioners (NARUC), 2002.

[53] Kevin R Keeney and James E Braun. Application of Building Precooling to Reduce Peak Cooling Requirements. ASHRAE Transactions, pages 463–469, 1997.

[54] Sila Kiliccote, Mary Ann Piette, Edward Koch, and Dan Hennage. Utilizing Automated Demand Response in Commercial Buildings As Non-Spinning Reserve Product for Ancillary Services Markets. In Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on, pages 4354–4360, 2011.

[55] Brendan Kirby. Load Response Fundamentally Matches Power System Reliability Requirements. In Power Engineering Society General Meeting, 2007. IEEE, pages 1–6, Tampa, FL, 2007.

[56] Moncef Krarti. An Overview of Artificial Intelligence-Based Methods for Building Energy Systems. Journal of Solar Energy Engineering, 125(3):331, 2003.

[57] Moncef Krarti, J F Kreider, D Cohen, and P Curtiss. Estimation of Energy Savings for Building Retrofits Using Neural Networks. Journal of Solar Energy Engineering, 120(3):211–216, 1998.

[58] John D Kueck, Brendan J Kirby, Moonis R Ally, and Keith Rice. Using Air Conditioning Load Response for Spinning Reserve. Technical report, Oak Ridge National Laboratory, 2009.

[59] P Lauret, F Miranville, H Boyer, F Garde, and Laetitia Adelard. Bayesian parameter estimation of convective heat transfer coefficients of a roof-mounted radiant barrier system. Journal of Solar Energy Engineering, 128:213 – 225, May 2006.

[60] Lawrence Berkeley Laboratory and Los Alamos National Laboratory. DOE-2 Engineers Manual (Version 2.1A). Technical report, US Department of Energy, 1982.

[61] Matthew Leach, Chad Lobato, Adam Hirsch, Shanti Pless, and Paul Torcellini. Technical Support Document : Strategies for 50% Energy Savings in Large Office Buildings. Technical Report September, National Renewable Energy Laboratory, Golden, CO, 2010.

[62] Kyoung-ho Lee and James E Braun. Development and Application of an Inverse Building Model for Demand Response in Small Commercial Buildings. In Proceedings of SimBuild, pages 1–12, Boulder, 2004.

[63] Kyoung-ho Lee and James E Braun. Development of Methods for Determining Demand-Limiting Setpoint Trajectories in Commercial Buildings Using Short-Term Data Analysis. In Second National IBPSA-USA Conference, pages 99–106, Cambridge, MA, 2006.

[64] Kyoung-Ho Lee and James E Braun. Development of methods for determining demand-limiting setpoint trajectories in buildings using short-term measurements. Building and Environment, 43(10):1769–1783, October 2008.

[65] Kyoung-ho Lee and James E Braun. Evaluation of methods for determining demand-limiting setpoint trajectories in buildings using short-term measurements. Building and Environment, 43(10):1755–1768, October 2008.

[66] Kyoung-ho Lee and James E Braun. Model-based demand-limiting control of building thermal mass. Building and Environment, 43:1633–1646, 2008.

[67] Uwe Ligges. tuneR: Analysis of music, 2011.

[68] Amory B Lovins. More Profit with Less Carbon. Scientific American, pages 74–82, September 2005.

[69] Kahlid Malik, Maurice Kugler, Milorad Kovacevic, William Orme, Eva Jespersen, and Sarantuya Mend. Human Development Report: Technical Notes. Technical report, United Nations Development Programme, New York, NY, 2013.

[70] Daniel M Martínez and Ben W Ebenhack. Understanding the role of energy consumption in human development through the use of saturation phenomena. Energy Policy, 36(4):1430–1435, April 2008.

[71] MATLAB. version 8.0.0.783 (R2012b). The MathWorks Inc., Massachusetts, 2012.

[72] Peter T May-Ostendorp, Gregor P Henze, Charles D Corbin, Balaji Rajagopalan, and Clemens Felsmann. Model-predictive control of mixed-mode buildings with rule extraction. Building and Environment, 46(2):428–437, February 2011.

[73] F B Morris, James E Braun, and Stephen Treado. Experimental and Simulated Performance of Optimal Control of Building Thermal Storage. ASHRAE Transactions, 100(1):402–414, 1994.

[74] C Neumann, D Jacob, S Burhenne, A Florita, E Burger, and F Schmidt. Model-based methods for fault detection and optimization in building operations (modellbasierte methoden für die fehlererkennung und optimierung im gebäudebetrieb). Technical report, Fraunhofer ISE, 2011.

[75] Simon J Olivieri. Evaluation of Commercial Building Demand Response Potential Using Optimal Short-Term Curtailment of HVAC Loads. Technical report, University of Colorado at Boulder - MS Thesis, 2011.

[76] Simon J Olivieri, Gregor P Henze, Chad D Corbin, and Michael J Brandemuehl. Evaluation of commercial building demand response potential using optimal short-term curtailment of heating, ventilation, and air-conditioning loads. Journal of Building Performance Simulation, 7(2):100–118, March 2014.

[77] Zheng O'Neill, Satish Narayanan, and Rohini Brahme. Model-Based Thermal Load Estimation in Buildings. In SimBuild, New York City, 2010.

[78] Andrew L Ott. Experience with PJM market operation, system design, and implementation. IEEE Transactions on Power Systems, 18(2):528–534, May 2003.

[79] Nadia S Ouedraogo. Energy consumption and human development: Evidence from a panel cointegration and error correction model. Energy, 63:28–41, December 2013.

[80] David Owen. The Efficiency Dilemma. New Yorker, 86(41):78–85, December 2010.

[81] N P Padhy. Unit CommitmentA Bibliographical Survey. IEEE Transactions on Power Systems, 19(2):1196–1205, May 2004.

[82] Alan D Pasternak. Global Energy Futures and Human Development : A Framework for Analysis. Technical Report October, Lawrence Livermore National Laboratory, Livermore, CA, 2000.

[83] PJM Interconnection. PJM Manual 28: Operating Agreement Accounting (Revision: 61). Technical report, Market Settlements Development Department, 2013.

[84] R Development Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, 2008. ISBN 3-900051-07-0.

[85] A Rabl and L K Norford. Peak Load Reduction By Preconditioning Buildings at Night. International Journal of Energy Research, 15:781–798, 1991.

[86] Ivan Rajšl, Perica Ilak, Marko Delimar, and Slavko Krajcar. Dispatch Method for Independently Owned Hydropower Plants in the Same River Flow. Energies, 5(12):3674–3690, September 2012.

[87] T Agami Reddy, Jason K Lukes, and Leslie K Norford. Benefits of Multi-Building Electric Load Aggregation: Actual and Simulation Case Studies. ASHRAE Transactions, 110, 2004.

[88] Yamina Saheb. Modernising Building Energy Codes. Technical report, International Energy Agency and the United Nations Development Programme, 2013.

[89] John E Seem. Modeling of Heat Transfer in Buildings. PhD thesis, Wisconsin-Madison, 1987.

[90] Devinderjit Sivia and John Skilling. Parameter estimation II. In Data Analysis: A Bayesian Tutorial, chapter 3, pages 35 – 77. Oxford University Press, 2006.

[91] Steve Sorrell, John Dimitropoulos, and Matt Sommerville. Empirical estimates of the direct rebound effect: A review. Energy Policy, 37(4):1356–1371, April 2009.

[92] Eric Sortomme. Combined Bidding of Regulation and Spinning Reserves for Unidirectional Vehicle-to-Grid. In Innovative Smart Grid Technologies (IGST), 2012 IEEE PES, pages 1–7, Washington, DC, 2011.

[93] Eric Sortomme and Mohamed A El-Sharkawi. Optimal Charging Strategies for Unidirectional Vehicle-to-Grid. IEEE Transactions on Smart Grid, 2(1):131–138, 2011.

[94] Eric Sortomme and Mohamed A El-Sharkawi. Optimal Scheduling of Vehicle-to-Grid Energy and Ancillary Services. IEEE Transactions on Smart Grid, 3(1):351–359, 2012.

[95] Benjamin K Sovacool. The intermittency of wind, solar, and renewable electricity generators: Technical barrier or rhetorical excuse? Utilities Policy, 17(3-4):288–296, September 2009.

[96] Ugur Soytas and Ramazan Sari. Energy consumption and GDP: causality relationship in G-7 countries and emerging markets. Energy Economics, 25(1):33–37, January 2003.

[97] Richard K Strand, Drury B Crawley, Curtis O Pedersen, Richard J Liesen, Linda K Lawrie, F C Winkelmann, W F Buhl, Y J Huang, and Daniel E Fisher. EnergyPlus : A New-Generation Energy Analysis and Load Calculation Engine for Building Design. In Proceedings of the ACSA Technology Conference, volume 2, Cambridge, Massachusetts, 2000. NREL.

[98] The World Bank. Energy use (kg of oil equivalent per capita), 2011.

[99] Paul A Torcellini, Michael Deru, Brent T Griffith, Kyle Benne, Mark Halverson, David Winiarski, and Drury B Crawley. DOE Commercial Building Benchmark Models. In ACEEE Summer Study on Energy Efficiency in Buildings, 2008.

[100] United Nations Development Programme. Table 1: Human Development Index and Its Components, 2013.

[101] Shengwei Wang and Xinhua Xu. Parameter estimation of internal thermal mass of building dynamic models using genetic algorithm. Energy Conversion and Management, 47:1927–1941, 2006.

[102] Shengwei Wang and Xinhua Xu. Simplified building model for transient thermal performance estimation using GA-based parameter identification. International Journal of Thermal Sciences, 45:419–432, 2006.

[103] B S Warr and R U Ayres. Evidence of causality between the quantity and quality of energy consumption and economic growth. Energy, 35(4):1688–1693, April 2010.

[104] Michael Wetter. Simulation-Based Building Energy Optimization. PhD thesis, University of California-Berkeley, 2004.

[105] Michael Wetter. A Modelica-based Model Library for Building Energy and Control Systems. In 11th International Building Performance Simulation Association Conference - Building Simulation, Glasgow, 2009.

[106] Michael Wetter. Modelica-based Modeling and Simulation to Support Research and Development in Building Energy and Control Systems. Journal of Building Performance Simulation, 2(2):143–161, 2009.

[107] Jiznxun Wu, T Agami Reddy, and David E Claridge. Statistical Modeling of Daily Energy Consumption in Commercial Buildings Using Multiple Regression and Principal Component Analysis. In Proceedings of the Eighth Symposium on Improving Building Systems in Hot and Humid Climates, Dallas, 1992.

[108] Peng Xu, Philip Haves, Mary Ann Piette, and James E Braun. Peak Demand Reduction from Pre-Cooling with Zone Temperature Reset in an Office Building. In 2004 ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, CA, 2004.

[109] Rongxin Yin, Peng Xu, Mary Ann Piette, and Sila Kiliccote. Study on Auto-DR and pre-cooling of commercial buildings with thermal mass in California. Energy and Buildings, 42(7):967–975, July 2010.

[110] Victor M Zavala. Real-Time Optimization Strategies for Building Systems. Industrial & Engineering Chemistry Research, 52(9):3137–3150, March 2013.

[111] Peng Zhao, Gregor P. Henze, Sandro Plamp, and Vincent J. Cushing. Evaluation of commercial building HVAC systems as frequency regulation providers. Energy and Buildings, 67:225–235, December 2013.

[112] M Zhou, Z Yan, Y X Ni, G Li, and Y Nie. Electricity price forecasting with confidence-interval estimation through an extended ARIMA approach. IEE Proceedings - Generation, Transmission and Distribution, 153(2):187–195, March 2006.

[113] Qiang Zhou, Shengwei Wang, Xinhua Xu, and Fu Xiao. A grey-box model of next-day building thermal load prediction for energy-efficient control. International Journal of Energy Research, pages 1418–1431, 2008.

[114] Na Zhu, Shengwei Wang, Xinhua Xu, and Zhenjun Ma. A simplified dynamic model of building structures integrated with shape-stabilized phase change materials. International Journal of Thermal Sciences, 49(9):1722–1731, 2010.

Appendix A

Portfolio 2: Thermal Mass Analysis

It was desired to analyze the impact of varying building thermal mass on the synergy observed within the portfolio. The portfolio from Case 2, consisting of 20 medium office buildings and one large office, was used to perform the simulation investigation. The thermal mass characteristics of the inverse gray-box RC models were modified by changing the capacitance assigned to the zone air node, effectively increasing or decreasing the available internal thermal storage.

A.1 Initial Testing and Experiment Design

As a first step, the NSU control strategies were simulated for mass levels of 0.5, 1, and 2 times the original mass to observe the model behavior. Figure A.1 highlights the NSU simulation results for the medium office building. At hour ending 4 AM the buildings start at virtually the same zone temperature. $M_{0.5}$ and M_1 are able to achieve the desired occupied zone temperature setpoint of 23.89 °C within one hour of start up, while M_2 requires a bit longer. The floating response of each building is observed after hour ending 17, and shows a slower temperature rise for the more massive buildings as expected. The corresponding electric consumption is shown in the bottom panel. At start up, the more massive buildings experience a larger load to reach the desired setpoint, as expected. Similar results are observed in Figure A.2 for the large office building; and, overall, the response appears to capture the desired characteristics of various building mass levels.

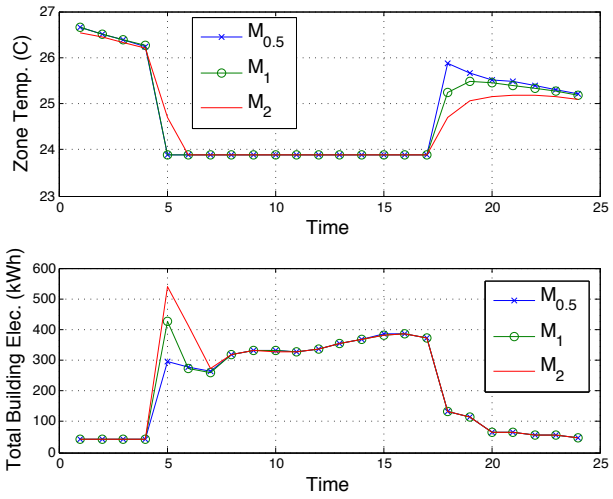

Figure A.1: NSU simulation of medium office with 0.5x, 1x, and 2x internal mass.

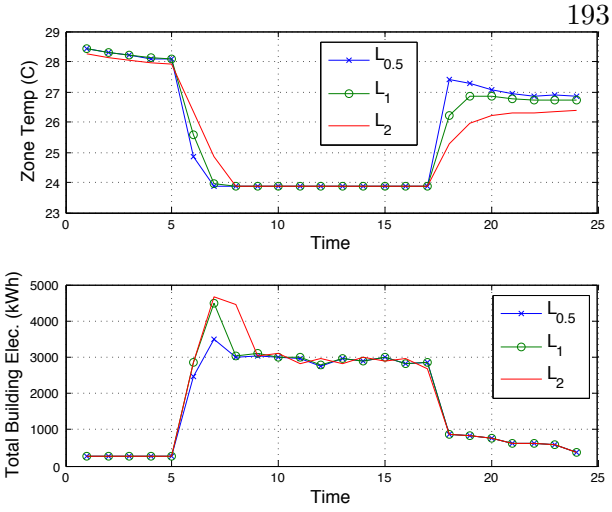

Figure A.2: NSU simulation of large office with 0.5x, 1x, and 2x internal mass.

The previous portfolio optimizations were constructed based on preliminary optimization results that determined the minimum peak achievable by each building. Setting the target demand too low forces all buildings into maximum peak reducing strategies, while setting it too high eliminates the joint objective of managing peak demand. Either case results in virtually no synergy. Preliminary individual optimizations were performed for each building and mass level to determine the minimum peaks achievable, since changing thermal mass characteristics can alter the peak reducing capability of a building. This step was necessary to frame a problem where all buildings (and mass levels) can achieve the desired demand reductions, allowing for fair comparison of synergy between the variations in thermal mass.

Figure A.3 shows the individual optimization results for the medium office building with three different mass levels. The target demand was set to a low limit to evaluate the maximum achievable peak reductions. The bottom panel shows that the OPT strategies for the three mass levels were able to achieve very similar demand reductions, suggesting that halving the original mass does not significantly deteriorate the buildings load shifting abilities. The zone temperatures in the top panel show that the lower mass levels need to precool to lower temperatures to achieve the same demand reductions (i.e. greater ΔT can make up for less storage capacity).

Figure A.4 highlights the individual optimization results for the large office building with

three different mass levels. Comparing OPT electric profiles in the bottom panel shows that higher mass levels are able to achieve slightly lower on-peak electric demand. The thermal mass discharging strategies are similar among all mass levels, however, higher mass levels were able to keep warmer zone temperatures while charging during early morning hours.

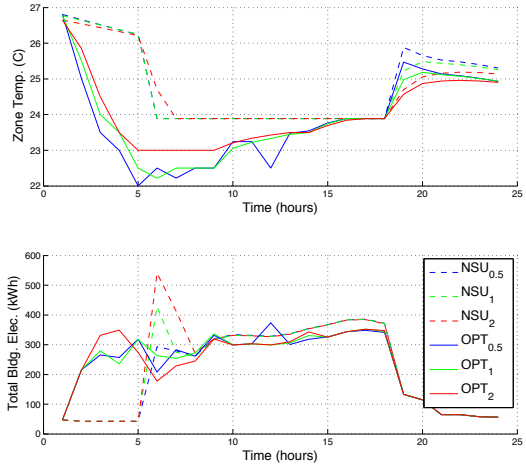

Figure A.3: Medium office results for various mass levels and low (i.e. unachievable) TDL.

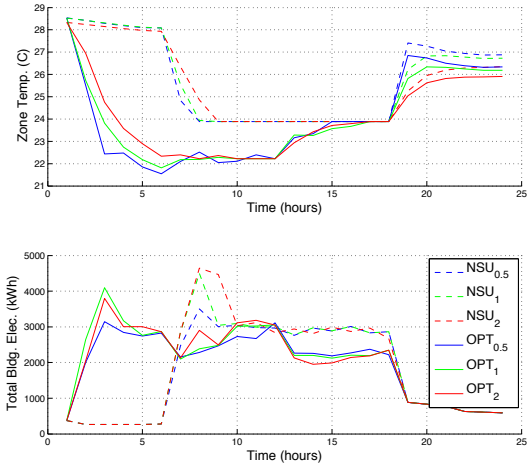

Figure A.4: Large office results for various mass levels and low (i.e. unachievable) TDL.

These results were used to revise the original demand limit settings used in the Case 2 portfolio. The portfolio target demand limit was slightly increased by 26 kW (0.26%) to 9915 kW so that demand reducing objectives could be achieved by all portfolio mass levels. The revised demand limits for scenarios a, b, and c are shown in Figure A.5. Scenario a still represents the case where nearly all peak reducing duties are given to the medium office buildings in the individual optimizations. Scenario b represents the case where portfolio demand reductions are split between the two building types. Scenario c represents the case where the large office building is given nearly all the peak reducing responsibility in the individual optimizations.

Three mass levels were explored for each building, which required nine portfolio optimizations to investigate all mass combinations. Exploring scenarios a, b, and c for each mass level required nine individual optimizations for each building, totaling 18 individual optimizations. The simulation experiment design and target demand limits are highlighted in Table A.1.

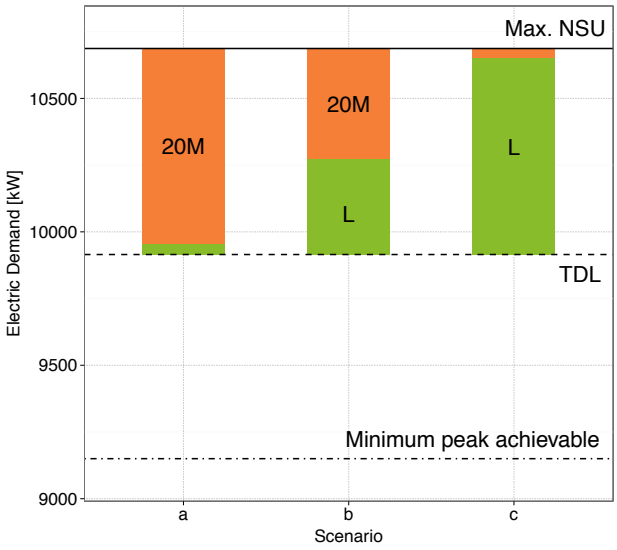

Figure A.5: Revised portfolio scenario target demand limits.

Table A.1: Optimization scenarios and target demand limits.

#	Case	Bldg	Scenario a	Scenario b	Scenario c	Portfolio	Unit
1	$M_{0.5}L_{0.5}$	20M L	6940 2975	7260 2655	7640 2275	9915	kW
2	$M_{0.5}L_1$	20M L	6940 2975	7260 2655	7640 2275	9915	kW
3	$M_{0.5}L_2$	20M L	6940 2975	7260 2655	7640 2275	9915	kW
4	$M_1L_{0.5}$	20M L	6940 2975	7260 2655	7640 2275	9915	kW
5	M_1L_1	20M L	6940 2975	7260 2655	7640 2275	9915	kW
6	M_1L_2	20M L	6940 2975	7260 2655	7640 2275	9915	kW
7	$M_2L_{0.5}$	20M L	6940 2975	7260 2655	7640 2275	9915	kW
8	M_2L_1	20M L	6940 2975	7260 2655	7640 2275	9915	kW
9	M_2L_2	20M L	6940 2975	7260 2655	7640 2275	9915	kW

A.2 Portfolio Results

The results of the 9 portfolio optimizations are highlighted in Figure A.6. Comparing the OPT and NSU cases for the medium and large office buildings shows that the optimal solution was to split demand reducing duties between both buildings for all combinations of thermal mass.

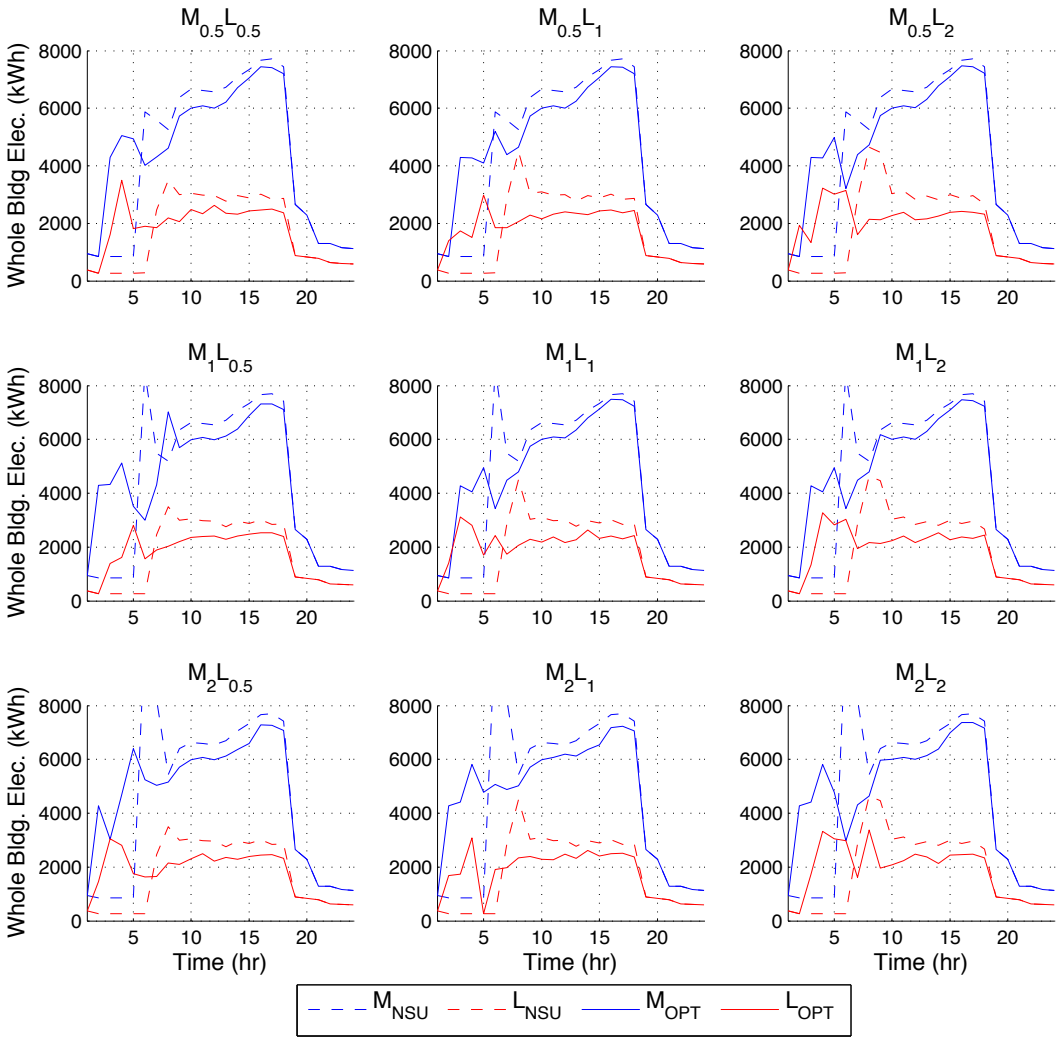

Figure A.6: Portfolio 2 demand profiles for variations in building thermal mass.

Plotting the results for all mass levels on the same axes highlights the slight variations in demand reduction for each building depending on the mass level. In Figures A.7 and A.8, the medium office mass level is represented by line type while large office mass level is represented by

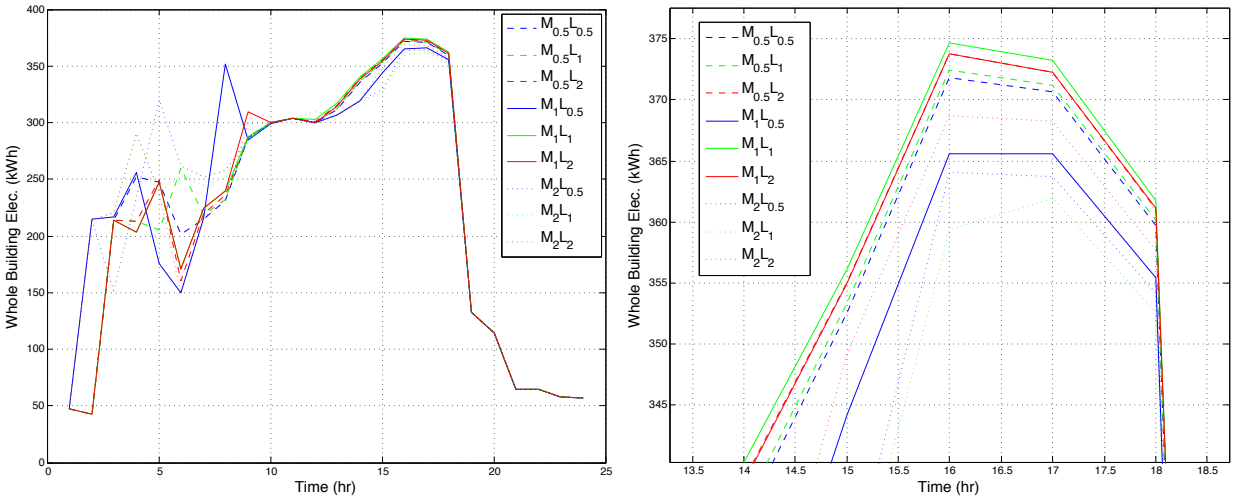

(a) Entire planning horizon (i.e. 24 hours). (b) Close-up of peak period.

Figure A.7: Medium office optimal electric profile for all mass variations.

line color. Although exceptions were observed, it appears that the medium office tended to make greater demand reductions when it had high mass (i.e. dotted lines are typically lower than dashed and solid lines in Figure A.7b), which seems logical.

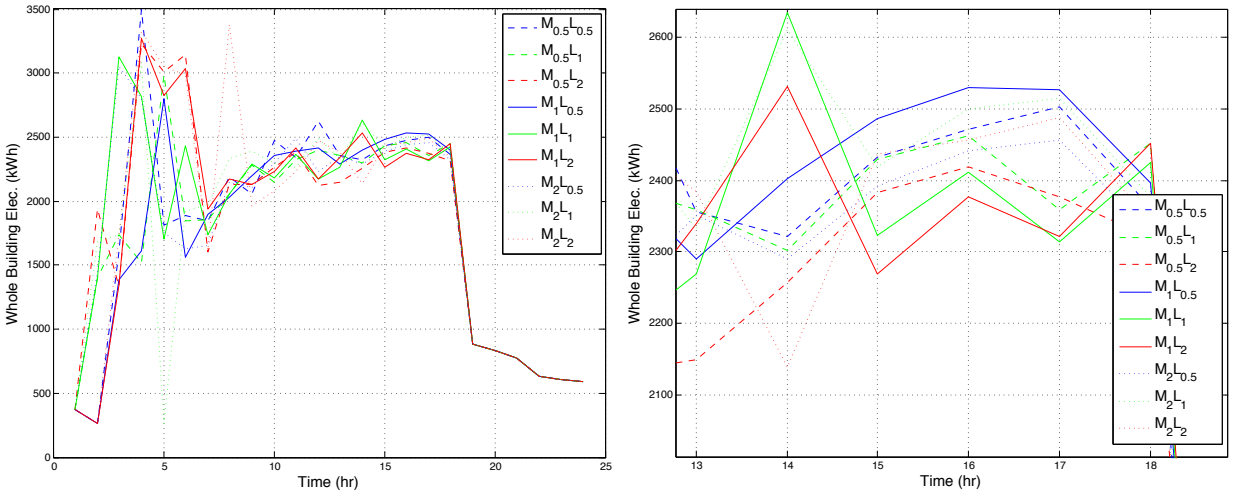

(a) Entire planning horizon (i.e. 24 hours). (b) Close-up of peak period.

Figure A.8: Large office optimal electric profile for all mass variations.

Similar results were observed in Figure A.8b, where the large office tended to contribute greater

demand reductions at higher mass levels (i.e. red lines are typically lower than green and blue lines).

A.3 Evaluation of Synergy

Synergy was evaluated by comparing the portfolio optimizations to the aggregated individual optimizations. The results are presented in terms of percent difference between the portfolio and individual results, where negative values indicate the portfolio achieved a lower value. Table A.2 shows the percent difference for optimizer total cost for scenarios a, b, and c. To distinguish patterns between scenarios cell color is scaled across all scenarios. As expected, no synergy was observed for scenario b, since this scenario is nearly identical to the optimal solution of the portfolios.

Table A.2: Percent difference in total optimizer cost.

Scenario a

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	-1.33	-1.54	-1.47
M_1	-0.81	-0.77	-1.12
M_2	0.03	0.29	0.14

Scenario b

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	0.22	-0.20	0.05
M_1	0.30	0.14	-0.04
M_2	0.52	0.57	0.60

Scenario c

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	-1.68	-1.40	0.01
M_1	-1.61	-1.09	-0.09
M_2	-1.54	-0.80	0.40

Scenario a results show that the highest amounts of synergy occured for the $M_{0.5}$ cases and tended to decrease as the medium office mass was increased. This result seems logical since scenario a is the case where the medium office was tasked with performing the majority of load reductions in the individual optimizations. Performing large load reductions in the medium office building with little thermal mass was less efficient, thus the $M_{0.5}$ cases could benefit the most from splitting demand reducing duties with the large office building. As mass was added to the medium office buildings they became more efficient at load shifting and the benefit of portfolio optimization decreased. No synergy was observed in scenario a for the cases with the medium office mass doubled (M_2).

Scenaro c results show that the highest amounts of synergy occurred for the $L_{0.5}$ cases and

tended to decrease as the large office mass was increased. Scenario c represents the case where the large office is tasked with the majority of demand reductions in the individual optimizations. Performing large load reductions with the low mass large office building was less efficient, thus greater opportunity for savings existed through splitting demand reductions with the medium office buildings. As the large office mass was increased it became more efficient at thermal storage and virtually no synergy was observed for the high mass case (i.e. L_2).

A similar relationship was observed in the overall energy consumption, suggesting that the majority of savings were related to energy savings rather than price arbitrage. It is also interesting to note that it appears that the energy savings tended to increase in scenario a as mass was added to the large office buildings. This may highlight the load shifting improvements of the higher mass large office building within the portfolio optimization. The opposite relationship was observed for scenario c, however, with energy savings tending to decrease as mass was added to the medium office buildings. This is largely attributable to the fact that increasing the mass of the medium office building created significant additional start-up loads to reach the desired occupied setpoint, resulting in higher energy consumption for higher mass levels.

Table A.3: Percent difference in energy consumption.

Scenario a

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	-4.25	-4.61	-4.47
M_1	-3.43	-3.64	-4.73
M_2	-0.26	-0.57	-1.39

Scenario b

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	-0.91	-1.32	-1.32
M_1	0.30	-1.04	-2.29
M_2	0.52	0.89	-0.07

Scenario c

	$L_{0.5}$	L_1	L_2
$M_{0.5}$	-2.96	-2.58	-0.45
M_1	-2.46	-1.93	-1.07
M_2	-1.85	-1.42	-0.26

Overall, the results of the simulation study suggest that opportunities for synergy may be greater when low mass building are present within the portfolio, since they can potentially benefit from offloading some (or all) of their demand reducing duties to other buildings. Optimizing portfolios consisting entirely of high thermal mass buildings resulted in no observed synergy, since

they can individually meet the desired demand reductions in a relatively efficient manner. The maximum synergy from a total cost perspective was observed in the $M_{0.5}L_{0.5}$ case for scenario c, which resulted in the portfolio optimization saving 1.68% (or $112) over the aggregated individual optimizations. Although the total cost savings may not be entirely compelling, almost 5% energy savings was observed for several cases in scenario a, which may be more significant.

Appendix B

HVAC Validation

B.1 Packaged RTU Validation

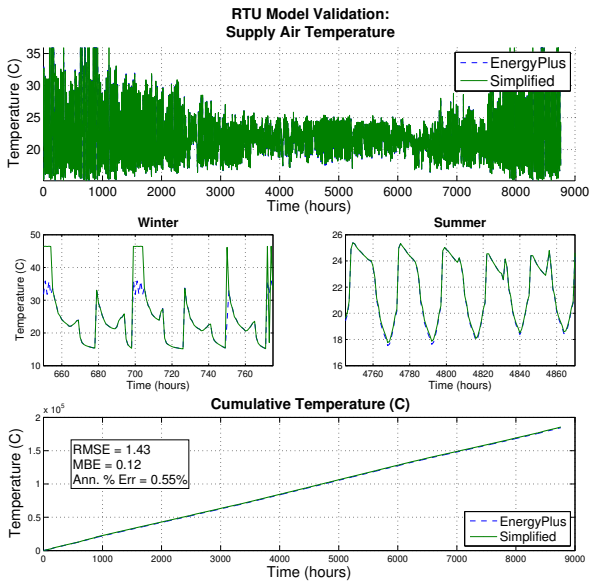

Figure B.1: Validation of packaged RTU supply air dry bulb temperature.

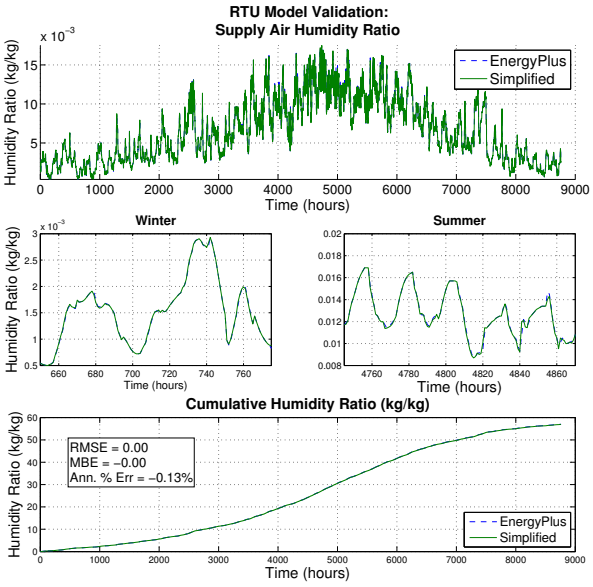

Figure B.2: Validation of packaged RTU supply air humidity ratio.

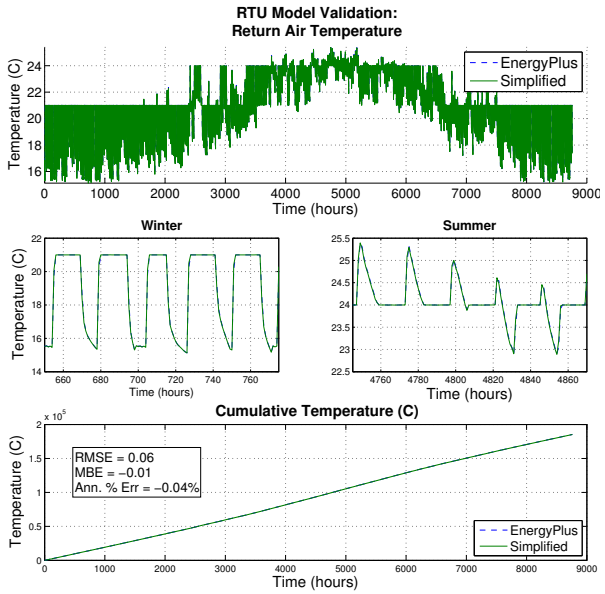

Figure B.3: Validation of packaged RTU return air dry bulb temperature.

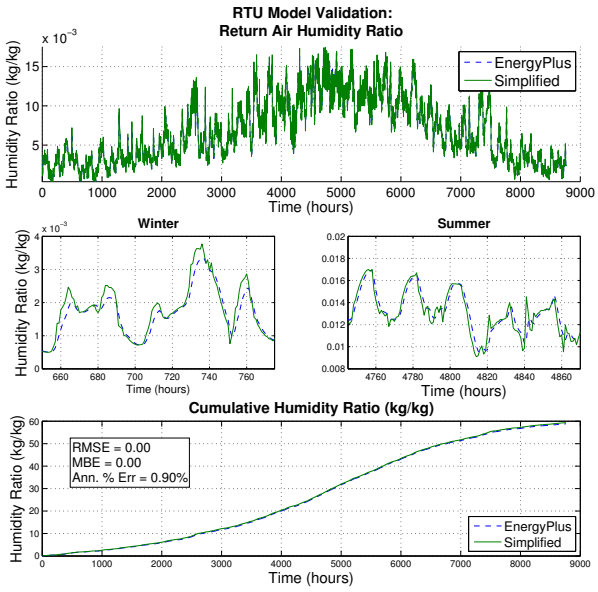

Figure B.4: Validation of packaged RTU return air humidity ratio.

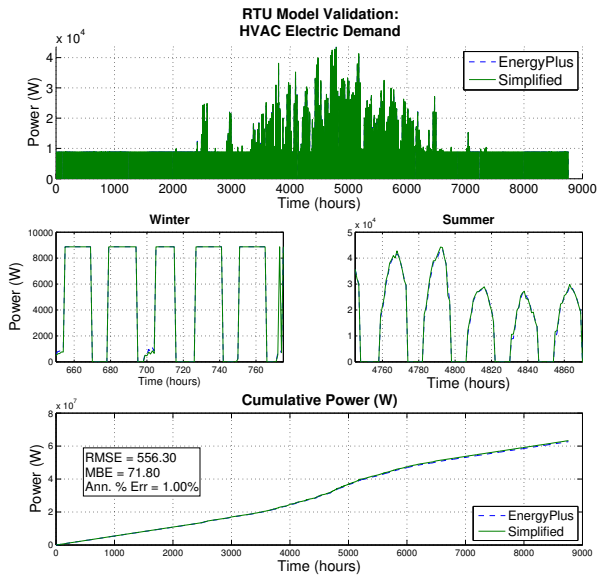

Figure B.5: Validation of packaged RTU electric demand.

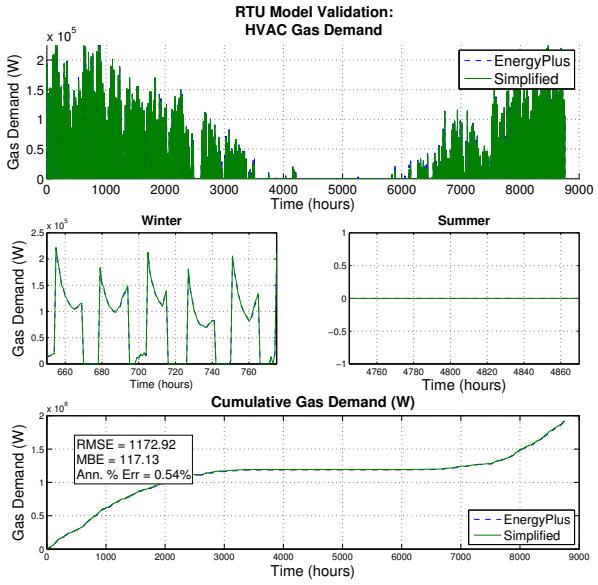

Figure B.6: Validation of packaged RTU gas demand.

B.2 Packaged VAV Validation

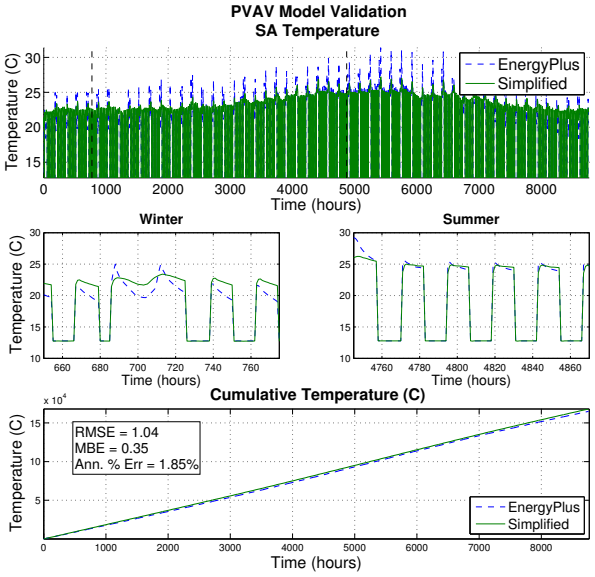

Figure B.7: Validation of packaged VAV supply air dry bulb temperature.

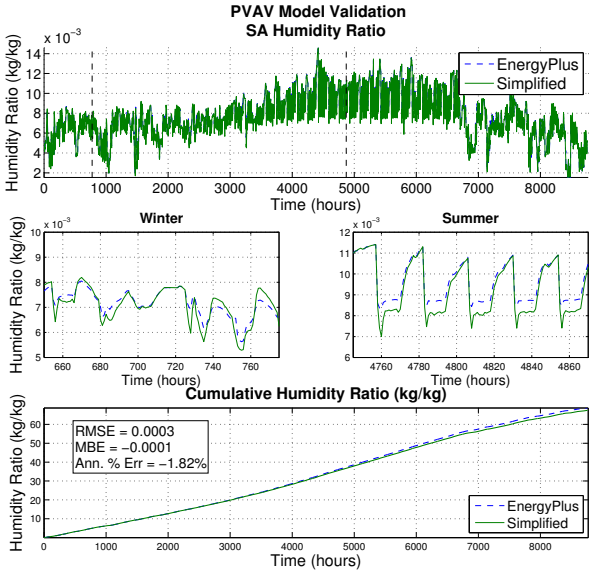

Figure B.8: Validation of packaged VAV supply air humidity ratio.

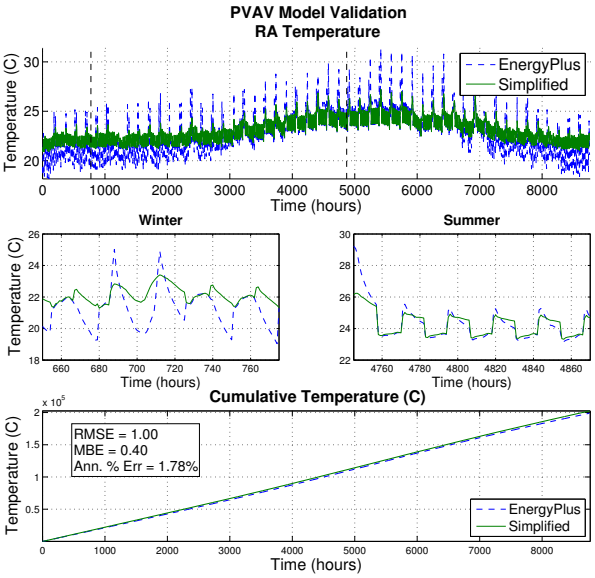

Figure B.9: Validation of packaged VAV return air dry bulb temperature.

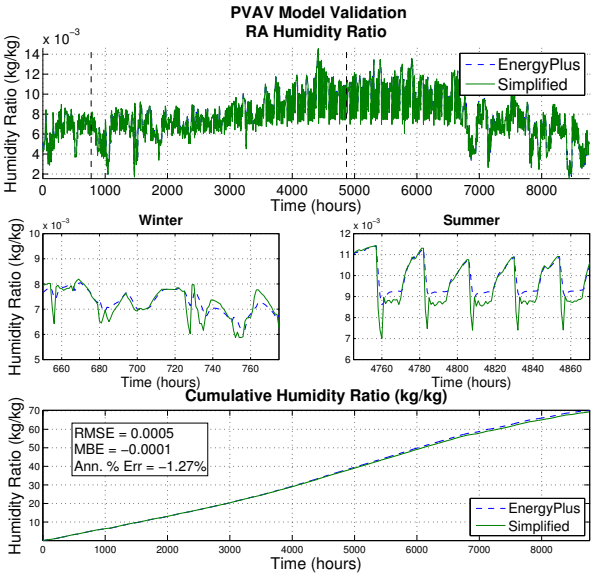

Figure B.10: Validation of packaged VAV return air humidity ratio.

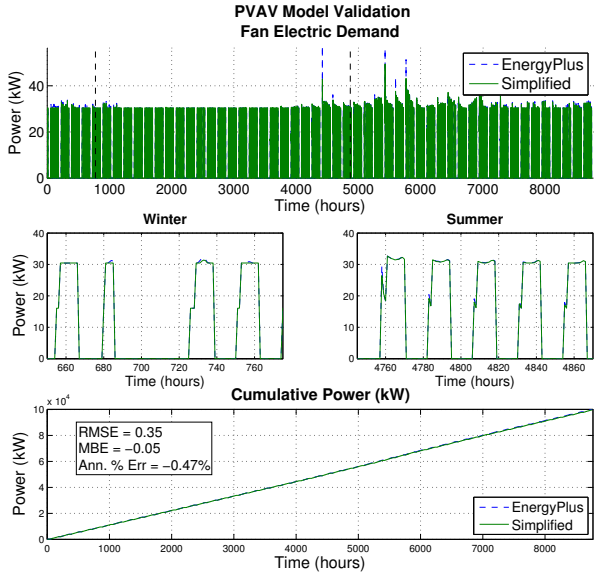

Figure B.11: Validation of packaged VAV fan power.

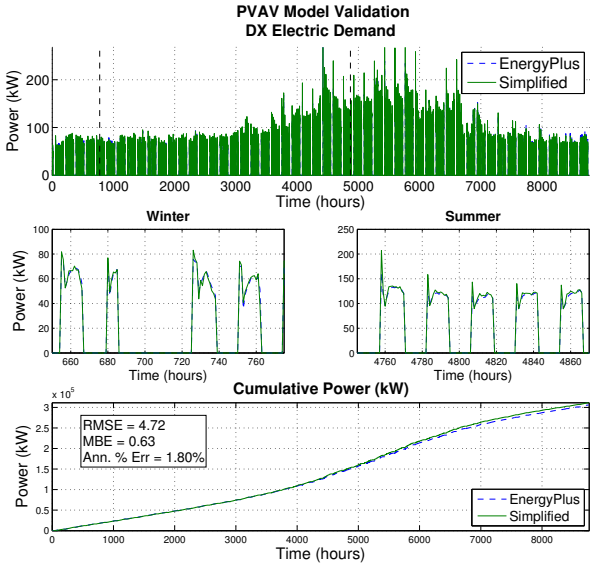

Figure B.12: Validation of packaged VAV DX coil condenser power.

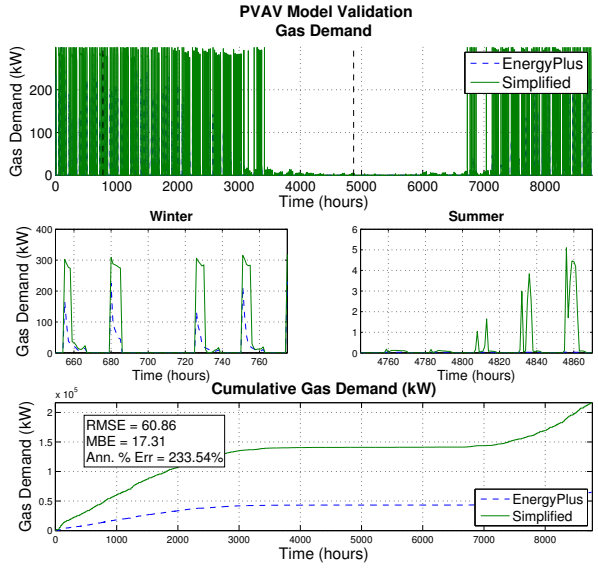

Figure B.13: Validation of packaged VAV gas demand.

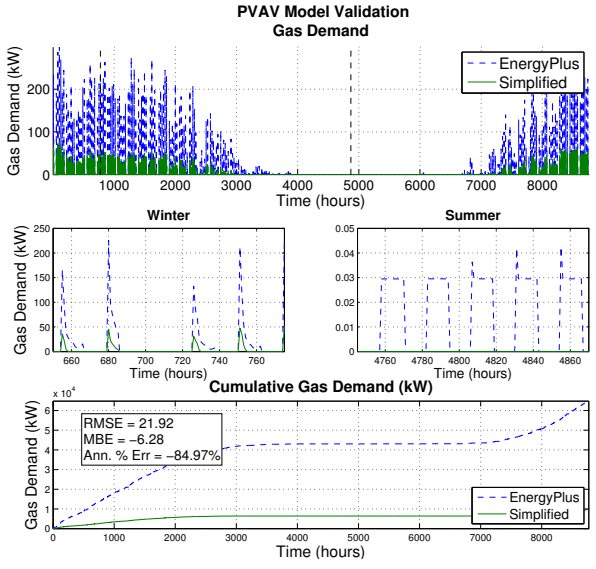

Figure B.14: Validation of packaged VAV gas demand (no vent).

B.3 Chiller Model Validation

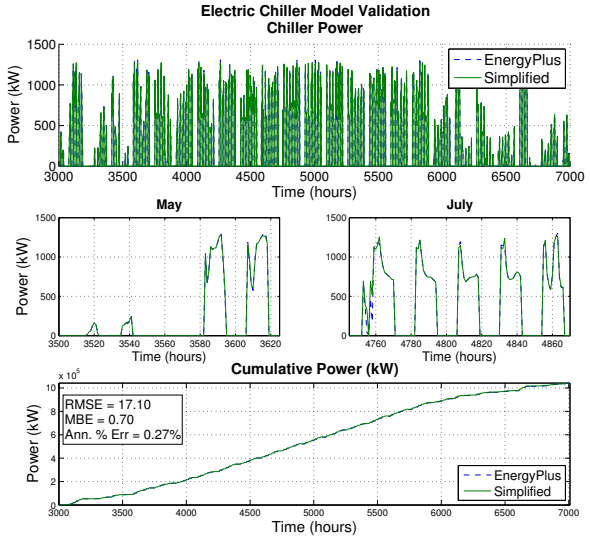

Figure B.15: Validation of chiller model power.

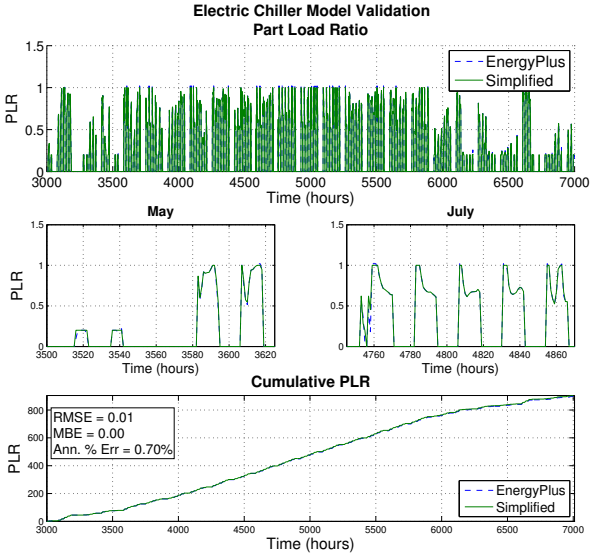

Figure B.16: Validation of chiller model PLR.

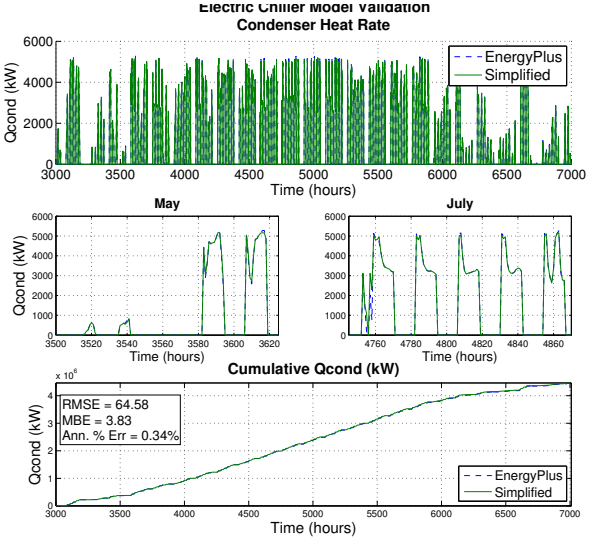

Figure B.17: Validation of chiller model condenser heat rate.

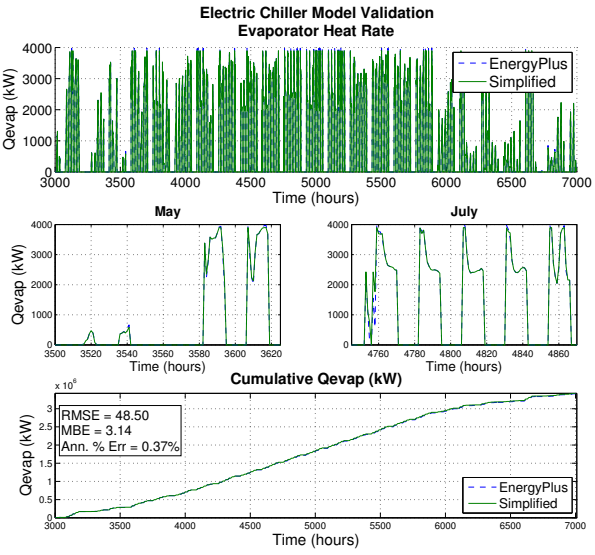

Figure B.18: Validation of chiller evaporator heat rejection.

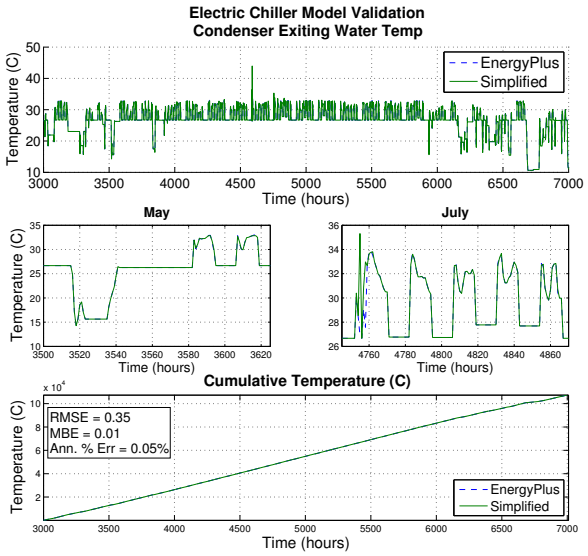

Figure B.19: Validation of chiller model condenser leaving water temperature.

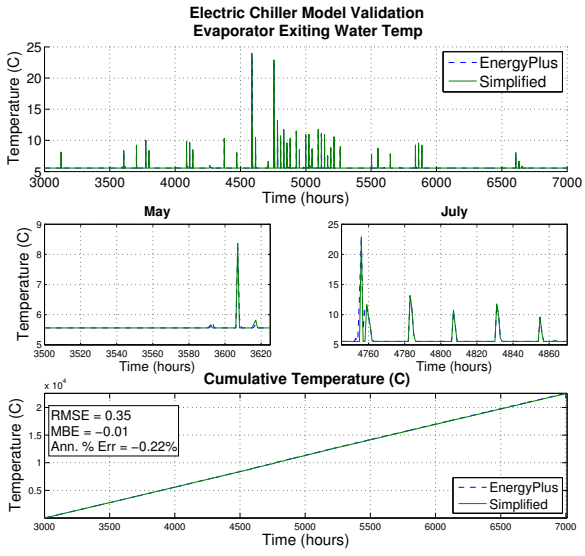

Figure B.20: Validation of chiller evaporator leaving water temperature.

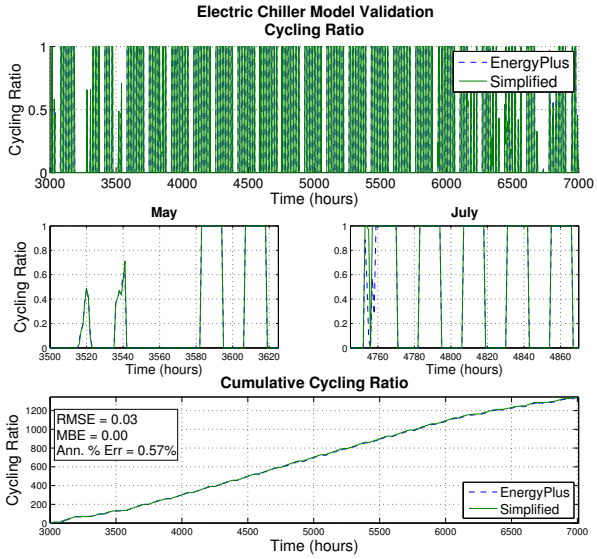

Figure B.21: Validation of chiller cycling ratio.

B.4 Chilled Water VAV System Validation

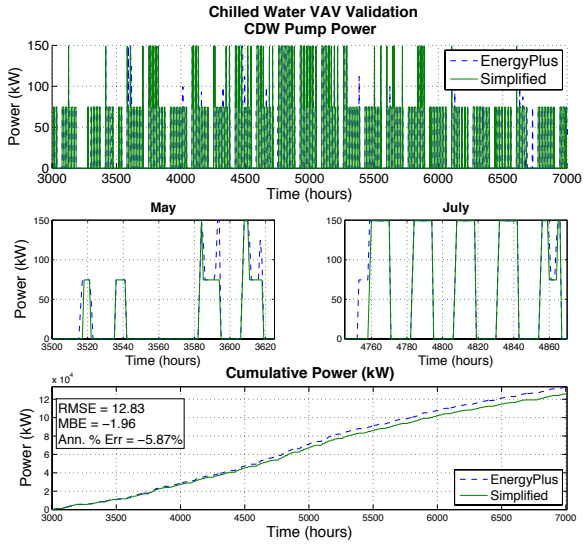

Figure B.22: Validation of condenser water pump power.

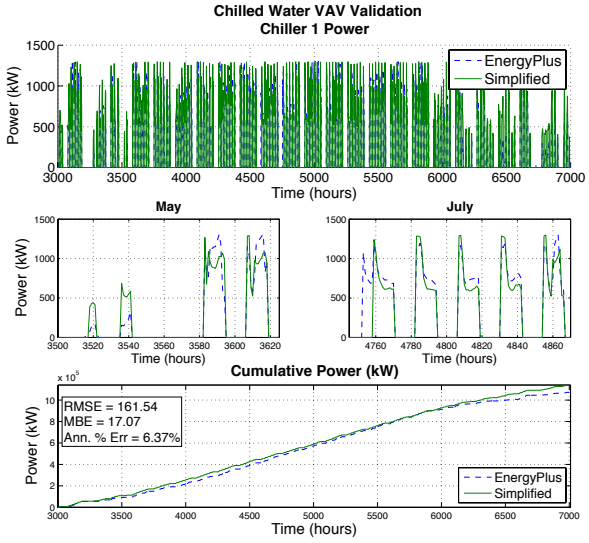

Figure B.23: Validation of chiller 1 power.

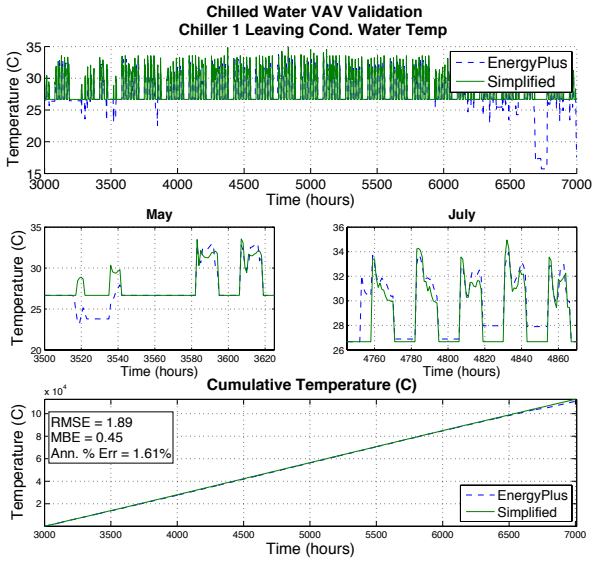

Figure B.24: Validation of chiller 1 condenser water outlet temperature.

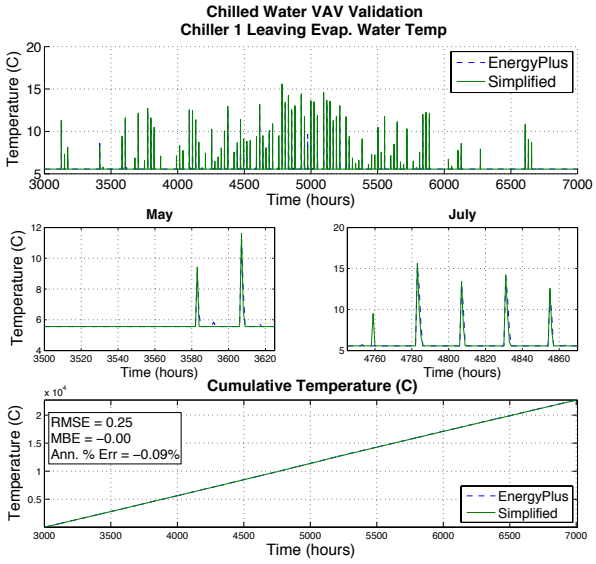

Figure B.25: Validation of chiller 1 evaporator water outlet temperature.

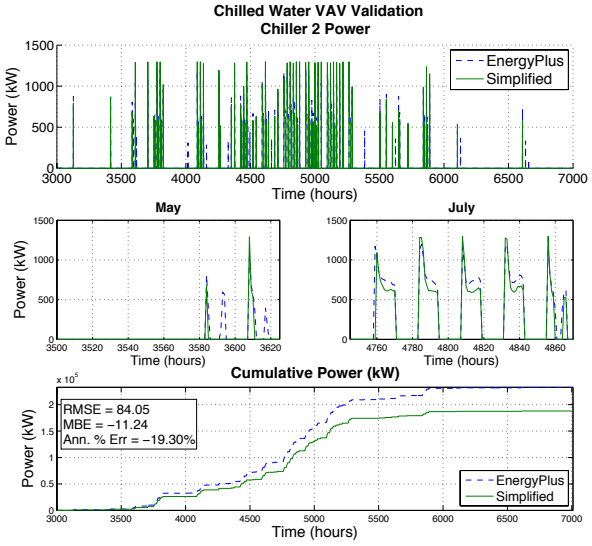

Figure B.26: Validation of chiller 2 power.

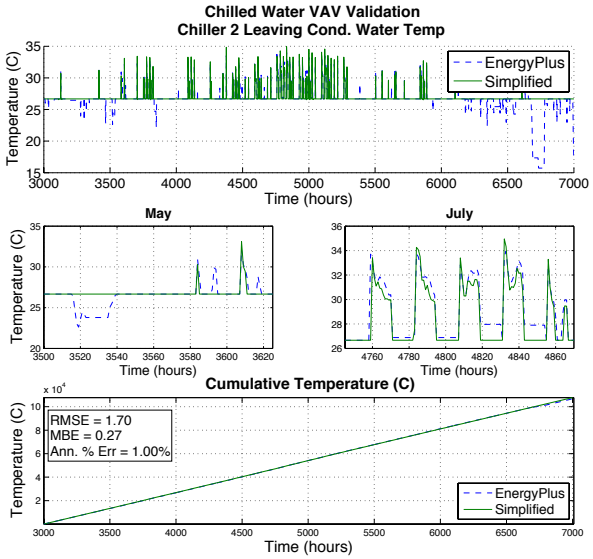

Figure B.27: Validation of chiller 2 condenser water outlet temperature

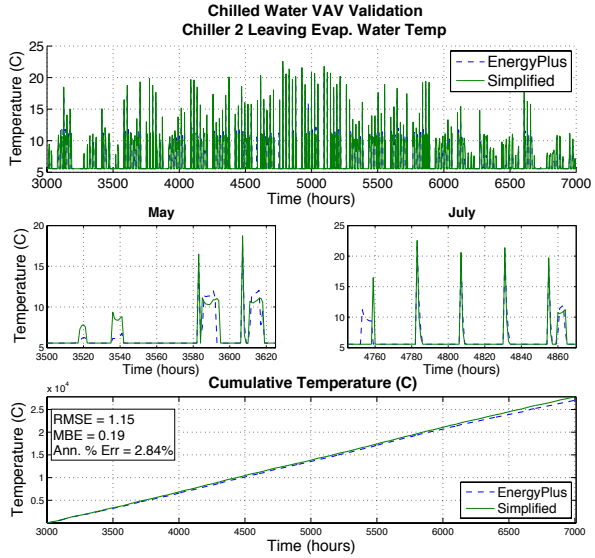

Figure B.28: Validation of chiller 2 evaporator water outlet temperature.

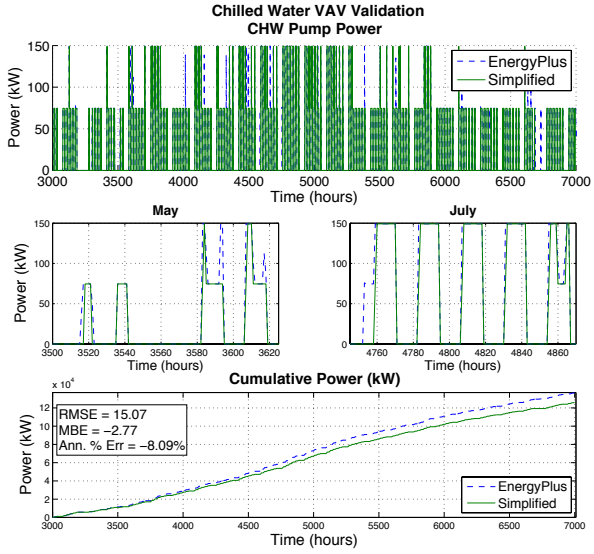

Figure B.29: Validation of chilled water pump power.

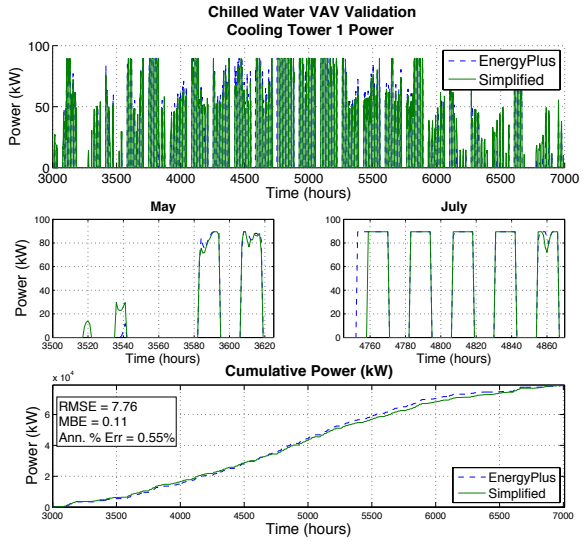

Figure B.30: Validation of cooling tower 1 power.

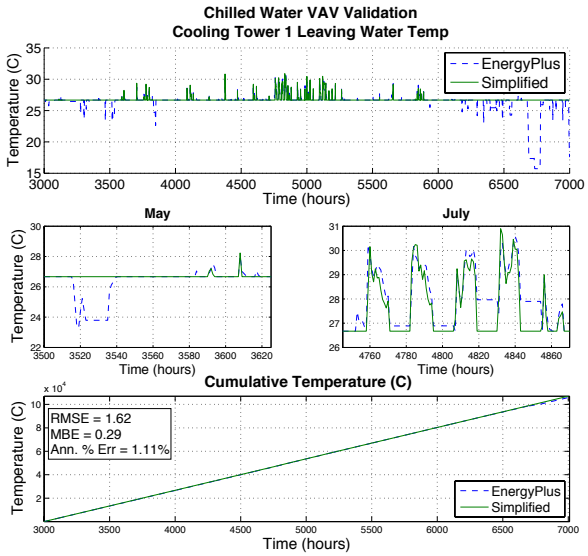

Figure B.31: Validation of cooling tower 1 outlet water temperature.

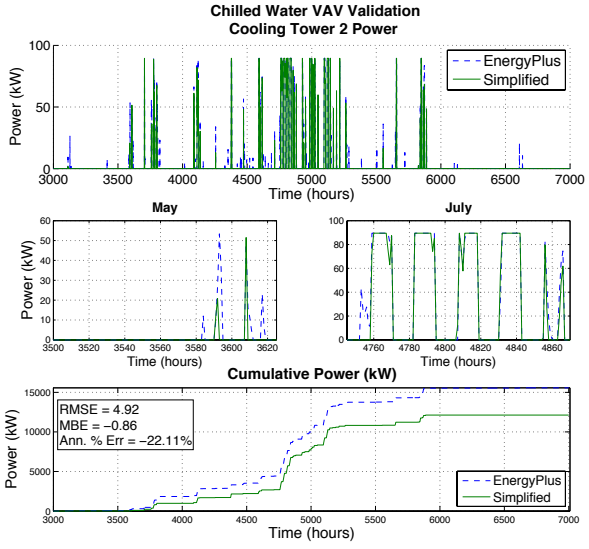

Figure B.32: Validation of cooling tower 2 power.

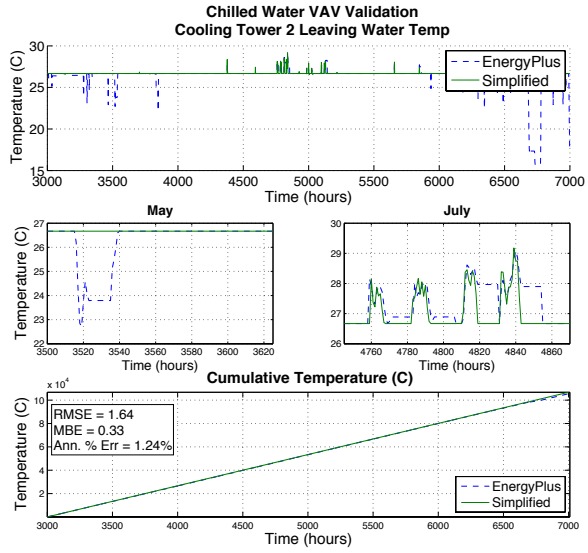

Figure B.33: Validation of cooling tower 2 outlet water temperature.

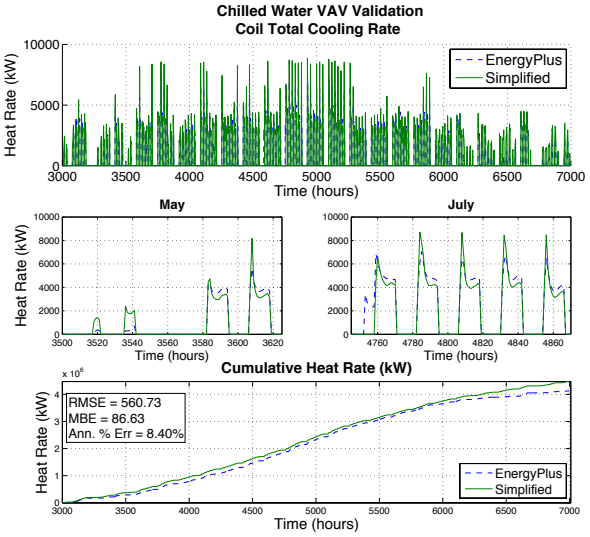

Figure B.34: Validation of cooling coil total cooling rate.

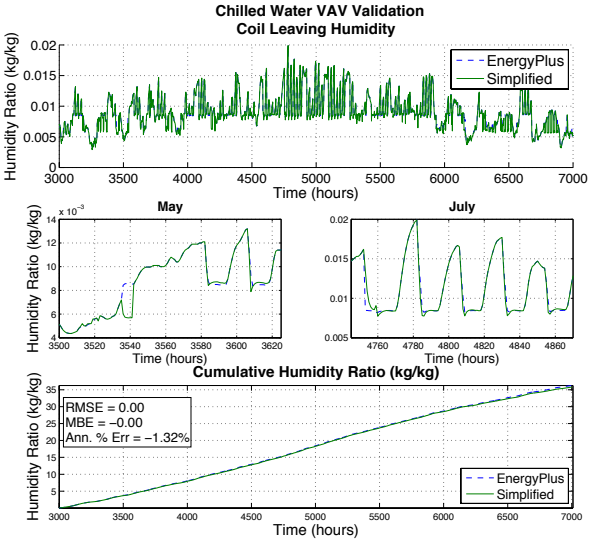

Figure B.35: Validation of cooling coil leaving humidity ratio.

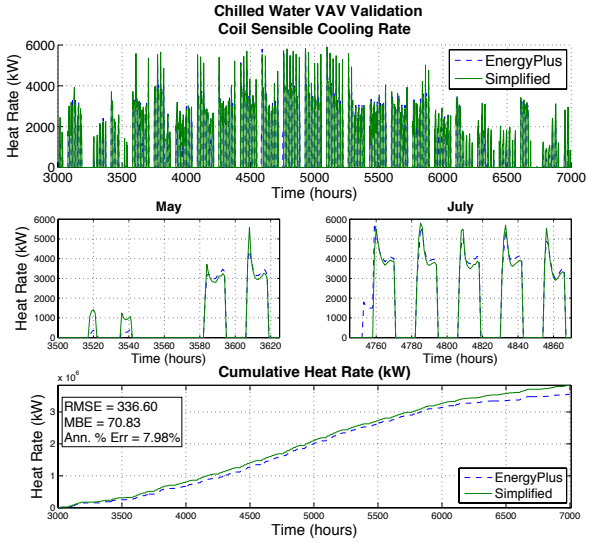

Figure B.36: Validation of cooling coil sensible cooling rate.

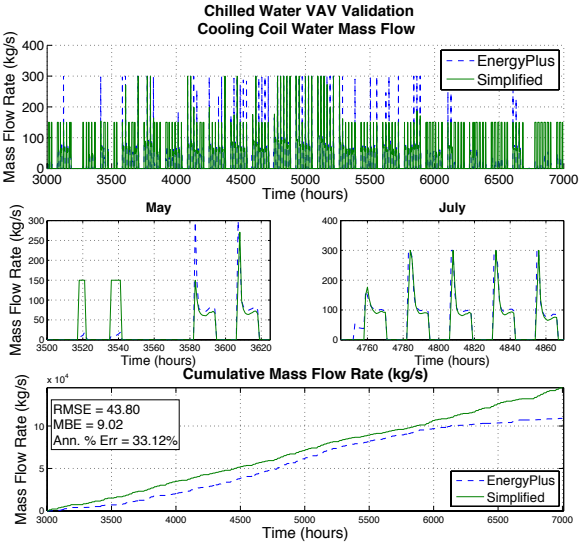

Figure B.37: Validation of cooling coil water mass flow rate.

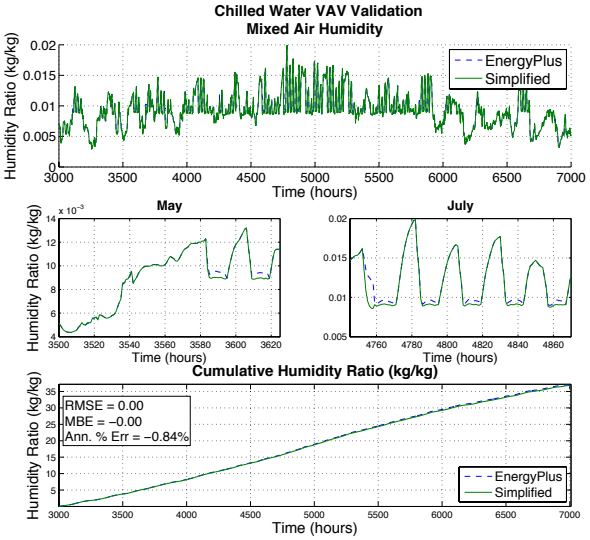

Figure B.38: Validation of mixed air humidity ratio.

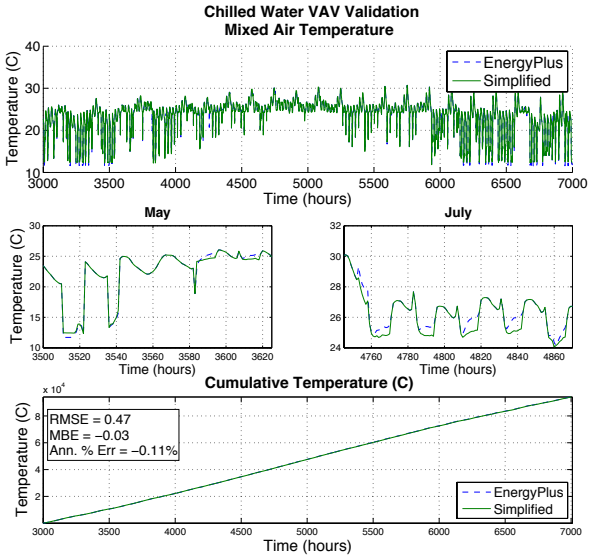

Figure B.39: Validation of mixed air temperature.

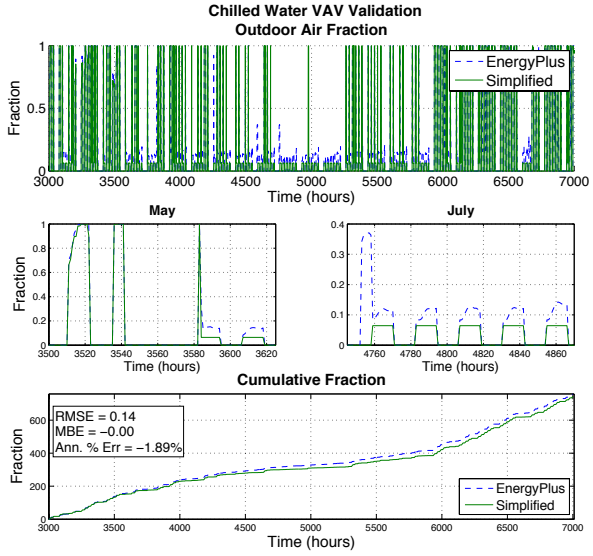

Figure B.40: Validation of outdoor air fraction.

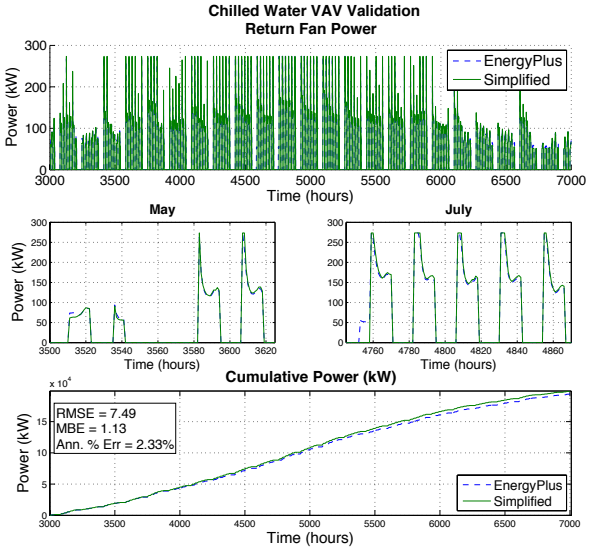

Figure B.41: Validation of return fan power.

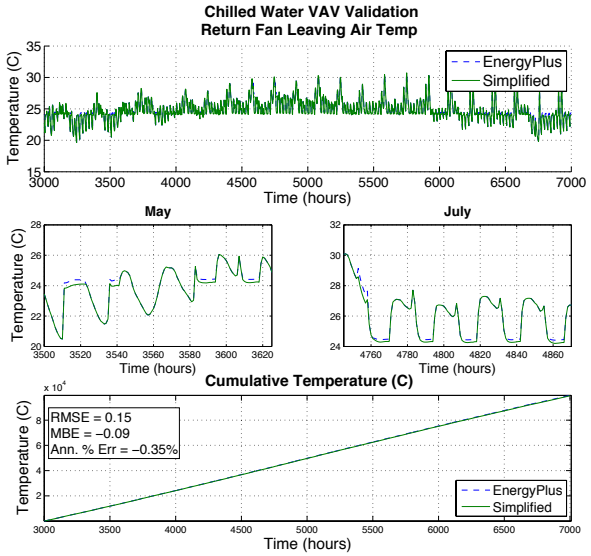

Figure B.42: Validation of return fan outlet air temperature.

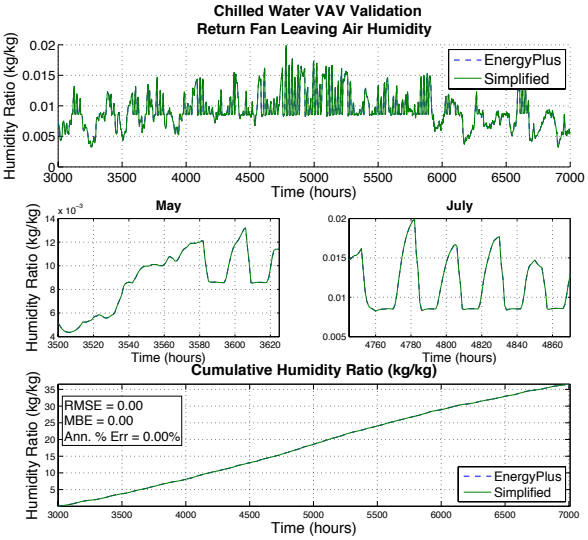

Figure B.43: Validation of return fan outlet humidity ratio.

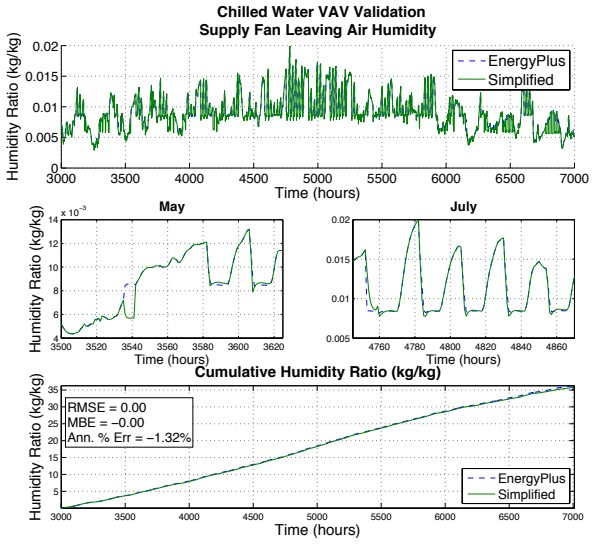

Figure B.44: Validation of supply fan leaving air humidity.

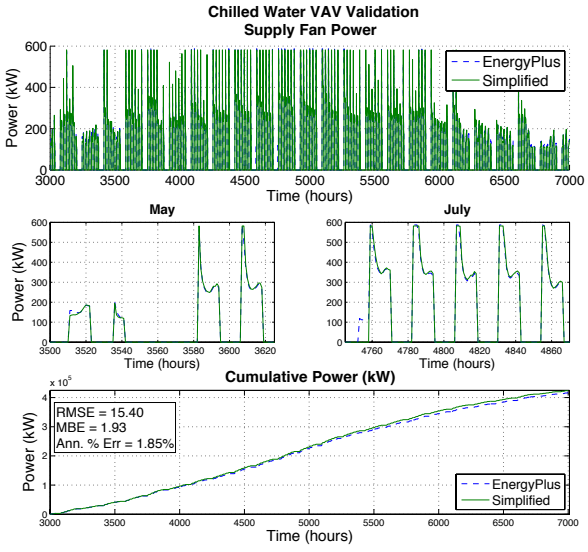

Figure B.45: Validation of supply fan power.

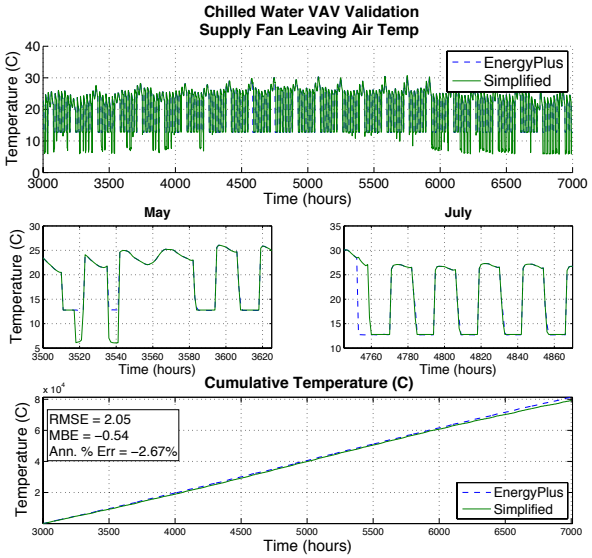

Figure B.46: Validation of supply fan leaving air temperature.

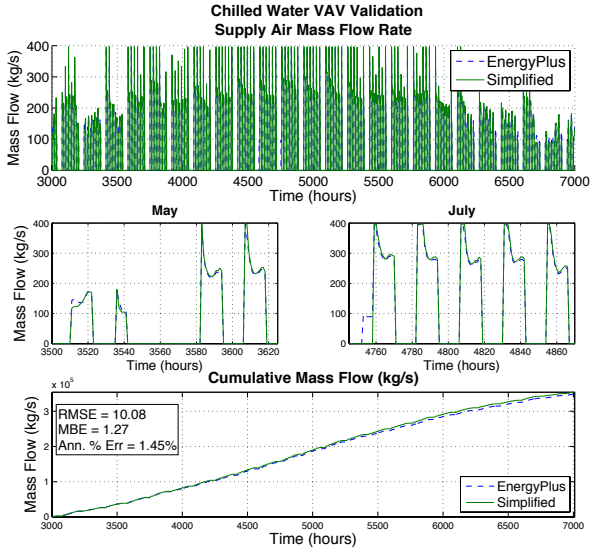

Figure B.47: Validation of supply air mass flow rate.

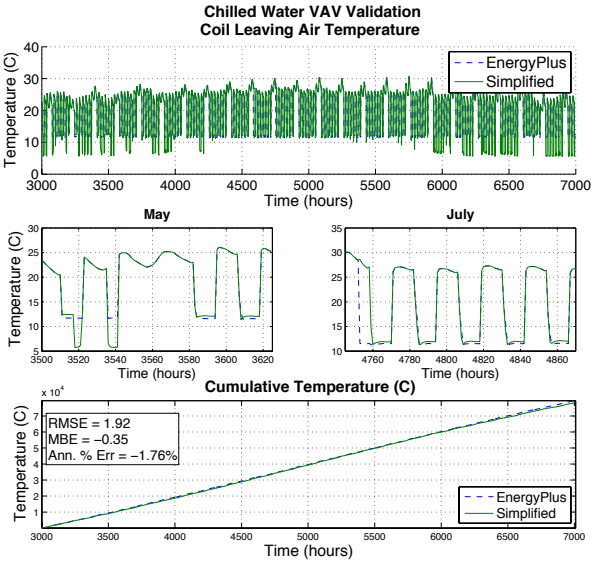

Figure B.48: Validation of cooling coil leaving air temperature.

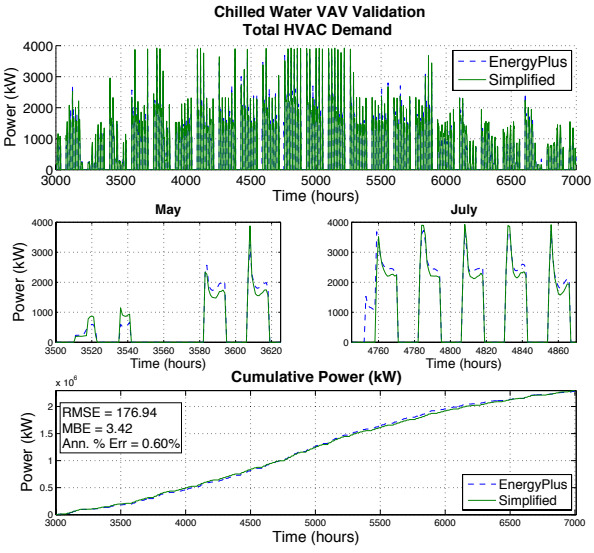

Figure B.49: Validation of total HVAC power.