MULTISCALE CONTROL OF ENERGY SYSTEMS

by

RANJEET KUMAR

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

Doctor of Philosophy
(Chemical Engineering)

at the
UNIVERSITY OF WISCONSIN-MADISON
2020

Date of final oral examination: May 11, 2020

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

Victor M. Zavala, Associate Professor, Chemical and Biological Engineering
Christos T. Maravelias, Professor, Chemical and Biological Engineering
Reid C. Van Lehn, Assistant Professor, Chemical and Biological Engineering
James Luedtke, Professor, Industrial and Systems Engineering

To my family.

ACKNOWLEDGMENTS

My journey in graduate school at the University of Wisconsin-Madison has been filled with numerous experiences and I am thankful to all the people who have been part of this journey. First of all, I would like to thank Prof. Victor Zavala with my deepest gratitude for mentoring me through my time in graduate school. Since my first meeting with him, I have always had immense respect and admiration for him for his expertise in the research area, his skills of communicating research, his patience, and his personality of being in good spirits at all times. His guidance and support, in both academic and non-academic matters, have been vital in keeping me motivated to constantly pursue my research with enthusiasm, and learning to develop critical thinking and innovative solutions to research questions. I consider myself very fortunate to have had the opportunity to work with him as one of his first Ph.D. students.

I would like to thank Prof. Christos Maravelias, Prof. Reid Van Lehn, and Prof. James Luedtke for on my dissertation committee. I am grateful to Prof. Maravelias for being my co-advisor for the first three years of my Ph.D. during which discussions with him provided great insights and useful feedback. Prof. Van Lehn also helped me with the perspective of researchers from outside the field of process systems engineering through which I learned to communicate my research to a wider audience. Prof. Luedtke's class on stochastic programming in Fall 2016 which was immensely helpful for me in my research. I am very grateful to Prof. Luedtke for his teaching and discussions during our course project which helped me learn the concepts that formed the foundation of my research.

During my time as a Ph.D. student at UW-Madison, I have had the chance to learn from some of the best researchers in the areas of process systems engineering and chemical engineering. In addition to my thesis committee members, I would like to express my gratitude to Professors James Rawlings, James Dumesic, Daniel Klingenberg, Steve Wright, and Alberto Del Pia for being my teachers here. I also want to thank all current and former members of Zavalab that provided me with an excellent peer group since the beginning. Apoorva Sampat and Jordan Jalving started their graduate studies in the group with me and have been great friends and colleagues. Special thanks to Jordan for helping me create beautiful animations for stochastic and deterministic MPC. Alex Dowling, Sungho Shin, Yicheng Hu, Yankai Cao, and all group members and visiting students who joined later have all been great friends, and I have enjoyed numerous discussions on various topics.

My graduate research would not have been possible without the constant support and feedback from our collaborators from Johnson Controls – Kirk Drees, Michael Wenzel,

Mohammad ElBsat, Matthew Ellis, and Michael Risbeck. I would also like to thank Robert Turney and Kirk Drees for providing me an opportunity for a summer internship at Johnson Controls in 2016 that helped me develop research questions relevant to the industrial setup. I am confident that the findings from my research will be useful to Johnson Controls for developing products in the future. I am also thankful to Victor Saucedo, Remya Pushpangatha Kurup, and Kawa Chiu of Genentech for the opportunity to work on stochastic programming for supply chain in pharmaceutical setting during my summer internship in 2018. The two industrial internships have been invaluable experiences for me.

I am quite fortunate to have made great friends, both in and outside of chemical engineering, during my time in Madison, which made my journey in graduate school even more enjoyable. It is not possible to write the names of all the friends here, but I would like to mention Saurabh Bhandari, Sandy Chen, Ashwin Shekar, and Arpit Agarwal in addition to Apoorva and Jordan, who have been a constant source of inspiration and support. Also, Atharva Kelkar and Akash Deep who became my neighbours in my last 2 years made my life in Madison a lot of fun. The numerous people who became friends in these years are all extremely valuable to me. During my graduate studies, I also volunteered as a part of Asha for Education - Madison Chapter, a non-profit organization that focuses on supporting the education of underprivileged children in India. This cause is very close to me because being from a small place in Bihar, the Indian state with the least literacy rate, I have experienced the impacts of lack of education closely. Through Asha for Education's Madison Chapter, I met several motivated and enthusiastic volunteers who were graduate students and people from working community, and in the four years, we disbursed over $40,000 to three projects in India that would benefit over 20,000 kids. I made long-lasting friendships as a volunteer at Asha for Education, and I appreciate their support to me.

My drive to pursue doctoral studies was a result of the excellent education that I received during my undergraduate studies at IIT Bombay that built strong foundations in me and made me confident to face challenges ahead in life. In particular, I want to thank Professors Ravindra Gudi, Sachin Patwardhan, and Hemant Nanavati who inspired and motivated me to pursue graduate studies. I will also note that my life-long friends from IIT Bombay have been a constant part of my life although living thousands of miles away.

It would have been impossible to achieve any of this without the support of my family. I have been one of the fortunate people from the place where I grew up to achieve the highest level of education. And, undoubtedly my parents are the main reason for this as they have provided all the support unconditionally in every aspect of my life. My elder brother has been a source of inspiration for me since an early age to keep working hard to achieve good results. My sister and younger brother have also been always there with their love and encouragement. Since getting married in February 2019, my wife has invariably supported me and kept me going through the last year of my Ph.D. I thank her for all the support and love she has provided me.

Ranjeet Kumar
Madison, WI
May 2020

CONTENTS

LIST OF FIGURES viii

LIST OF TABLES xiii

ABSTRACT xiv

1 INTRODUCTION 1
- 1.1 Motivation . . . 1
- 1.2 Outline . . . 3
- 1.3 Background . . . 5
  - 1.3.1 Energy Networks . . . 5
  - 1.3.2 Electricity Market Operations . . . 6
  - 1.3.3 Energy Storage in Electricity Markets . . . 10
  - 1.3.4 Model Predictive Control . . . 12
- 1.4 An Electricity Market Case Study . . . 13
  - 1.4.1 Modeling and Economic Analysis of Participation in a multiscale Energy Market in Deterministic Setting . . . 13
  - 1.4.2 Results and Economic Analysis in California ISO (CAISO) . . . 16
  - 1.4.3 Summary . . . 17

2 STOCHASTIC MODEL PREDICTIVE CONTROL FOR STATIONARY BATTERY SYSTEMS 18
- 2.1 Introduction . . . 19
- 2.2 Problem Formulation . . . 22
  - 2.2.1 Decision-Making Setting . . . 22
  - 2.2.2 Deterministic MPC Framework . . . 23
- 2.3 Stochastic MPC Framework . . . 30
  - 2.3.1 Objective Function . . . 31
  - 2.3.2 Constraints . . . 31
- 2.4 Benchmarking Stochastic MPC . . . 34
  - 2.4.1 Performance Metrics . . . 34
  - 2.4.2 Scenario Generation from Limited Historical Data . . . 36
  - 2.4.3 Computational Experiments . . . 39

2.5 Conclusions . . . 50

3 BENCHMARKING STOCHASTIC AND DETERMINISTIC MPC: A CASE STUDY IN STATIONARY BATTERY SYSTEMS **51**
3.1 Introduction . . . 51
3.2 Computational Framework . . . 54
3.2.1 Preliminaries . . . 54
3.2.2 Forecasting and Uncertainty Quantification . . . 57
3.2.3 Deterministic MPC . . . 58
3.2.4 Stochastic MPC . . . 61
3.2.5 Perfect Information MPC . . . 63
3.2.6 Handling Constraint Violations . . . 63
3.2.7 Benchmarking Procedure . . . 64
3.3 Assessment Results . . . 66
3.3.1 Closed-Loop Behavior . . . 67
3.3.2 Economic Performance . . . 70
3.3.3 FR Shortfall and Constraint Violations . . . 74
3.3.4 FR Buffer Tuning . . . 76
3.4 Conclusions . . . 79

4 STOCHASTIC MODEL PREDICTIVE CONTROL FOR CENTRAL HVAC PLANTS **80**
4.1 Introduction . . . 81
4.2 Computational Framework . . . 83
4.2.1 Decision-Making Setting . . . 84
4.2.2 Forecasting and Uncertainty Quantification . . . 87
4.2.3 Deterministic MPC . . . 88
4.2.4 Stochastic MPC . . . 94
4.2.5 Perfect Information MPC . . . 97
4.2.6 Handling Constraint Violations . . . 97
4.2.7 Benchmarking Procedure . . . 98
4.3 Benchmark Results . . . 100
4.3.1 Forecasting of Loads and Electricity Prices . . . 101
4.3.2 Closed-Loop Performance . . . 106
4.3.3 Economic Performance and Constraint Violations . . . 108
4.4 Conclusions . . . 114

5 HIERARCHICAL MPC SCHEMES FOR PERIODIC SYSTEMS USING STOCHASTIC PROGRAMMING **115**
5.1 Introduction . . . 116
5.2 Basic Definitions and Setting . . . 118
5.3 Hierarchical MPC Schemes . . . 123
5.3.1 Retroactive Optimization . . . 124
5.3.2 Incremental Cutting Plane Scheme . . . 125

5.3.3 Short-Term MPC Controller . . . 133
5.3.4 Extensions to Nonlinear Systems . . . 135
5.3.5 Stability Considerations . . . 136
5.4 Computational Experiments . . . 138
5.5 Conclusions . . . 144

6 A STOCHASTIC DUAL DYNAMIC PROGRAMMING FRAMEWORK FOR MULTISCALE MPC **145**
6.1 Introduction . . . 146
6.2 Deterministic Multistage Formulation . . . 149
6.3 Stochastic Multistage Formulation . . . 153
6.3.1 Stochastic Dual Dynamic Programming . . . 155
6.3.2 Extended MPC Formulations . . . 157
6.4 Case Study . . . 158
6.5 Conclusions . . . 163

7 DUAL DYNAMIC PROGRAMMING FOR MULTISCALE MIXED-INTEGER MPC **164**
7.1 Introduction . . . 164
7.2 Mixed-Integer MPC Formulation . . . 167
7.2.1 Dynamic Programming Representation . . . 168
7.2.2 Dual Dynamic Integer Programming . . . 170
7.2.3 Multiscale Mixed-integer MPC Formulations . . . 173
7.3 Computational Experiments . . . 175
7.3.1 Decision-Making Setting . . . 175
7.3.2 Mixed-Integer MPC Formulation for Central HVAC Plant . . . 179
7.3.3 Results . . . 182
7.4 Conclusions and Future Work . . . 191

8 CONCLUSIONS AND FUTURE DIRECTIONS **193**
8.1 Contributions . . . 194
8.1.1 Stochastic MPC Framework for Electricity Market Participation and Demand Charge Mitigation . . . 194
8.1.2 Computational Framework for Benchmarking Stochastic and Deterministic MPC . . . 194
8.1.3 A Retroactive Hierarchical MPC Scheme for Periodic Systems to Handle Long Horizons . . . 195
8.1.4 Dual Dynamic Programming Based Schemes for Long-Horizon MPC 195
8.2 Future research directions . . . 196
8.2.1 Extensions of the Stochastic MPC Framework to More Complex Market Conditions . . . 196
8.2.2 Future Directions on Benchmarking Framework for Stochastic MPC 196
8.2.3 Extensions of Retroactive Hierarchical MPC Schemes . . . 197
8.2.4 More Developments in Dual Dynamic Programming Schemes . . . 198

A NOMENCLATURE **199**
A.1 Nomenclature for Chapter 3 . . . 199
A.2 Nomenclature for Chapter 4 . . . 201
A.3 Nomenclature for Chapter 7 . . . 206

B HANDLING LONG HORIZONS IN MPC: A STOCHASTIC PROGRAMMING APPROACH **211**
B.1 Introduction . . . 211
B.2 Hierarchical MPC Scheme . . . 213
B.3 Stationary Battery Study . . . 217
B.3.1 Long-Term MPC Formulation . . . 218
B.3.2 Short-Term MPC Formulation . . . 221
B.3.3 Analysis for Imperfect Forecasts . . . 222
B.4 Results . . . 224
B.4.1 Perfect Forecasts . . . 226
B.4.2 Imperfect Forecasts . . . 226
B.5 Conclusions . . . 229

C A HIERARCHICAL MODEL PREDICTIVE CONTROL APPROACH FOR HANDLING DEMAND CHARGES USING BATTERY SYSTEMS **230**
C.1 Introduction . . . 230
C.2 Problem Formulation . . . 232
C.3 Stationary Battery Case Study . . . 236
C.3.1 Long-Term MPC Formulation . . . 237
C.3.2 Short-Term MPC Formulation . . . 240
C.4 Results . . . 241
C.4.1 Perfect Forecasts . . . 243
C.4.2 Imperfect Forecasts . . . 246
C.5 Conclusions . . . 247

BIBLIOGRAPHY **249**

LIST OF FIGURES

1.1 Sketch of energy hub in a utility network (Geidl and Andersson, 2007). 2
1.2 Energy prices for a day in the markets in California. The day-ahead market (DAM) prices vary at 1-hour timescale, the quarter-hourly market (QHM) prices are at 15-minute timescale, the real-time market (RTM) prices are at 5-minute resolution. 8
1.3 Illustration of the electricity market interactions with various energy technologies. 9
1.4 Peak shaving application of energy storage systems (Figure adapted from Oudalov et al. (2007b)). 11
1.5 Illustration of transactions of energy between the power grid and battery system. Quantities in parentheses denote the price for the corresponding transaction. 13
1.6 Illustration of the reserved energy storage for regulation capacity (given by constraints (1.4.5) and (1.4.6)). 14

2.1 FR signal (real-time and hourly average) from PJM. 23
2.2 Interactions between battery system, ISO, and utility. 24
2.3 Sketch of FR capacity and FR dispatch signal. 26
2.4 Sketch of stochastic MPC scheme. 33
2.5 Historical loads and realizations 39
2.6 Historical energy prices and realizations 40
2.7 Historical FR prices and realizations 41
2.8 Historical hourly FR signals and realizations 42
2.9 Distribution of total monthly costs (1000 realizations) 44
2.10 Distribution of peak demands over a month (1000 realizations) 45
2.11 Peak shaving performance for stochastic, deterministic, and perfect information MPC in a random load profile (7-day horizon, discounted demand) 46
2.12 Distributions of FR commitments (7-day horizon, discounted demand) 46
2.13 SOC policy in stochastic, deterministic and perfect information MPC (7-day horizon, discounted demand charge) 47

3.1 Combined battery and attached load system. 56
3.2 Sketch of computational framework. 66

3.3 Closed loop profile for deterministic MPC with $\rho = 0.1$. Black lines represent forecasts and model predictions. Blue lines represent actual realizations. For the control policies, red lines represent the actual implemented policy. For the state of charge, red lines represent the FR buffer. For the residual load, the green line represents the running peak R_t. 68
3.4 Closed loop profile for stochastic MPC with $\rho = 0$. The grey regions represent uncertainty forecasts with black representing mean forecasts. Blue lines represent realized observations or committed policies. For the control policies, red lines represent the actual implemented policy and usually overlap with committed policy. For the residual load, the green line represents the peak observed residual. 69
3.5 Comparison of FR shortfall for deterministic and stochastic MPC for a validation scenario. 70
3.6 Probability and cumulative distributions for total costs. 71
3.7 Probability and cumulative distributions for demand charges. 72
3.8 Probability and cumulative distributions for the value of the battery. 73
3.9 Probability and cumulative distributions for FR shortfall. 75
3.10 Probability and cumulative distributions for constraint violations. 75
3.11 Economic performance (total cost) of deterministic MPC with varying FR buffers. 78
3.12 Economic performance (value of battery) of deterministic MPC with varying FR buffers. 78

4.1 Schematic representation of central HVAC plant under study. 85
4.2 Historical electrical load of the campus. Red vertical lines denote end of each month. 102
4.3 Historical hot water load of the campus. Red vertical lines denote end of each month. 102
4.4 Historical chilled water load of the campus. Red vertical lines denote end of each month. 103
4.5 Historical electricity price data. Red vertical lines denote end of each month. . 103
4.6 A single instance of electrical load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time. . 104
4.7 A single instance of chilled water load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time. 105

4.8 A single instance of hot water load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time. . 105
4.9 A single instance of electricity price forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time. . 106
4.10 Closed-loop profile for deterministic MPC with $\beta = 0.1$. Black lines represent forecasts and model predictions. Blue lines represent actual realizations. For the control policies, red lines represent the actual implemented policy. For the state of charge, red horizontal lines represent the storage buffer. For the residual electrical load, the green line represents the running peak R_t. 107
4.11 Closed-loop profile for stochastic MPC with $\beta = 0$. The grey regions represent uncertainty forecasts with black representing mean forecasts. Blue lines represent realized observations or committed policies. For the control policies, red lines represent the actual implemented policy and usually overlap with committed policy. For the residual load, the green line represents the peak observed residual electricity demand. 108
4.12 Probability and cumulative distributions for total cost of central plant. 111
4.13 Probability and cumulative distributions for demand charges. 112
4.14 Probability and cumulative distributions for constraint violations (per 100 hours). 113
4.15 Economic performance (cost of central plant) of deterministic MPC with varying buffers. 114

5.1 Sketch of hierarchical scheme using cutting planes. 133
5.2 Energy price (top), FR price (middle), and load (bottom) data. 139
5.3 Optimal and optimal periodic policies. 141
5.4 Evolution of optimality gap. 141
5.5 Evolution of periodic SOC (top) and peak (bottom) targets obtained with cutting-plane scheme. 142
5.6 Evolution of periodic SOC targets and intra-period policies obtained with CP scheme (first seven periods). 143
5.7 Comparison of evolution of periodic SOC resulting from the retroactive approach and proactive MPC approach of Huang et al. (2011). 144

6.1 Full (top) and compact (bottom) scenario tree representations for multi-stage SP. 147
6.2 Illustration of battery system 159
6.3 Load scenarios for a week of operation. 160
6.4 Evolution of lower and upper bounds in SDDP. 163

7.1 Schematic representation of central HVAC plant under study. 177

7.2 Historical electrical load of the campus. Red vertical lines denote end of each month. . . . 181
7.3 Historical hot water load of the campus. Red vertical lines denote end of each month. . . . 181
7.4 Historical chilled water load of the campus. Red vertical lines denote end of each month. . . . 182
7.5 Historical electricity price data. Red vertical lines denote end of each month. . . . 182
7.6 **Top:** Evolution of upper and lower bounds over iterations for 168-hour horizon, continuous MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 168-hour horizon, continuous MPC. . . . 184
7.7 Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 168-hour horizon, continuous MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution. . . . 185
7.8 **Top:** Evolution of upper and lower bounds over iterations for 168-hour horizon, mixed-integer MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 168-hour horizon, mixed-integer MPC. . . . 186
7.9 Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 168-hour horizon, mixed-integer MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution. . . . 187
7.10 **Top:** Evolution of upper and lower bounds over iterations for 720-hour horizon, mixed-integer MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 720-hour horizon, mixed-integer MPC. . . . 189
7.11 Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 720-hour horizon, mixed-integer MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution. . . . 190

B.1 Hierarchical MPC scheme. . . . 217
B.2 Interactions battery, buildings, ISO, and utility. . . . 219
B.3 Market price data used for the case studies. . . . 223
B.4 Comparison of battery SOC (top), battery power (middle) and demand (bottom) policies. . . . 225
B.5 SOC policy without periodicity (top). Difference of SOC policies with and without periodicity constraints (right). . . . 227
B.6 Forecasted and realized load profiles. . . . 228

C.1 Hierarchical MPC scheme. 236
C.2 Interactions between battery, buildings, ISO, and utility. 237
C.3 Market and load data used for the case studies. 242
C.4 Comparison of SOC policies. 243
C.5 Comparison of battery power and building demand policies. 244
C.6 SOC policy without periodicity (left). Comparison of SOC policies with and without periodicity (right). 245
C.7 Forecasted and realized load profiles. 246

LIST OF TABLES

1.1 Revenues for different market participation schemes and operating mode combinations. Here, RTM includes both 15-minute and 5-minute timescales. 16

2.1 Expected costs for MPC schemes. TC= Total cost, DC=Demand charge, FR=frequency regulation, and E=Energy. 49
2.2 Value of stochastic solution (VSS) and value of perfect information (VPI). 49
2.3 Value of the battery under different MPC schemes. 49

3.1 Economic performance analysis for different MPC implementations. 74
3.2 Battery analysis using deterministic MPC with mean, minimum, and maximum forecasts for FR signal (with buffer of $\rho = 0.1$). 76
3.3 Expected costs for MPC schemes with varying FR buffers. 77
3.4 Expected costs and value of battery for deterministic MPC scheme with varying FR buffers. 77

4.1 Utility usage analysis for different MPC implementations. 109
4.2 Economic performance analysis for different MPC implementations (campus with central plant). 109
4.3 Economic performance analysis for different MPC implementations (central plant only). 110
4.4 Expected costs for deterministic MPC with varying buffers. 113

B.1 Comparison of cost items under perfect forecasts. 224
B.2 Comparison of cost items under imperfect forecasts. 229

C.1 Comparison of cost items under perfect forecasts. 244
C.2 Comparison of cost items under imperfect forecasts. 247

ABSTRACT

Energy networks are operated at multiple timescales (from hours to milliseconds) to ensure that supply and demands are matched in real time. Such coordination is achieved through hierarchical (multiscale) market transactions operated by the Independent System Operators (ISOs). There is a great interest in designing automation architectures to exploit the fast timescale electricity markets using energy storage systems. Energy storage systems like batteries provide dynamic flexibility to modulate electrical loads at various frequencies and aid the control of distribution and transmission networks. Generating revenue in electricity markets, however, is challenging because there is significant uncertainty on electricity prices and on revenue streams that can be generated from different markets (day-ahead, real-time, regulation).

A major accomplishment of this research is the development of multiscale stochastic model predictive control (MPC) framework that can be applied to systems such as energy storage systems like batteries or central HVAC plants for a large campus participating in the electricity markets and performing peak shaving applications. The stochastic MPC framework seeks to determine optimal energy allocation portfolios and buy/sell policies that maximize revenue from day-ahead and real-time electricity markets, and frequency regulation markets under uncertainties. The stochastic MPC controller incorporates a two-stage stochastic programming formulation that makes control decisions while taking into account uncertainty in predictions. This computational framework integrates forecasting, uncertainty quantification, and model predictive control (MPC) to benchmark the performance of deterministic and stochastic MPC. Through exhaustive closed-loop simulations in two case studies – a battery management case study and a central HVAC plant case

study, with real data from ISOs and a typical university campus, we illustrate that off-the-shelf deterministic MPC implementations suffer significant losses in performance and constraint violations due to their inability to handle disturbances that cannot be adequately represented by mean (most likely) forecasts.

We also developed a hierarchical model predictive control framework using stochastic programming to handle long horizon (or infinite horizon) problems for periodic systems. We show that if the state policy of an infinite-horizon problem is periodic, the problem can be cast as a retroactive stochastic program (SP) that progressively accumulates historical data to deliver the optimal periodic states. The retroactive problem can be seen as a high-level hierarchical layer that provides targets to guide a low-level MPC controller that operates over a short period at high time resolution. We derive a retroactive scheme tailored to linear systems by using cutting plane techniques and suggest strategies to handle nonlinear systems and to analyze stability properties.

Finally, to provide a scalable approach to handle complex MPC applications with uncertainties evolving over long time horizons and with fine time resolutions, we derived a stochastic dual dynamic programming (SDDP) framework from the perspective of MPC. Scalability is enabled by the use of a nested cutting-plane scheme, which uses forward and backward sweeps along the time horizon to adaptively construct and refine cost- to-go functions. We also leverage the nested decomposition scheme of SDDP to develop a dual dynamic programming scheme for long-horizon mixed-integer MPC problems.

1

INTRODUCTION

1.1 Motivation

Energy systems are becoming increasingly decentralized and exhibit new forms of coupling. In 2012, the US federal government increased the combined heat and power (CHP) capacity by 50%, from 80 GW to 120 GW, to promote industrial energy efficiency, which could save manufacturers as much as $100 billion in energy costs over the next decade (The White House, 2012). CHP and combined cooling, heat, and power (CCHP) systems are multi-generation systems that are positioned close to the demand points to provide multiple energy carriers (electricity, steam, hot/chilled water) to buildings and energy-intensive processes (Chicco and Mancarella, 2009). University campuses, manufacturing facilities, and urban utility networks are examples of facilities with such configurations (illustrated in Figure 1.1). These configurations provide opportunities for a wide range of applications of heat recovery and energy storage technologies. With heat recovery systems, energy efficiencies of up to 75% can be achieved in contrast to around 45% achieved by power plants located at significant distances from the demand locations. Energy storage systems in the form of electrochemical and thermal energy provide dynamic flexibility to the coupled energy networks and can be used to modulate the electrical loads at var-

ious frequencies (Geidl and Andersson, 2007). Energy storage can also aid the control of the distribution and transmission networks (Geidl and Andersson, 2007). This flexibility is key to manage large amounts of intermittent and non-dispatchable renewable power and ensure energy quality in the electric networks.

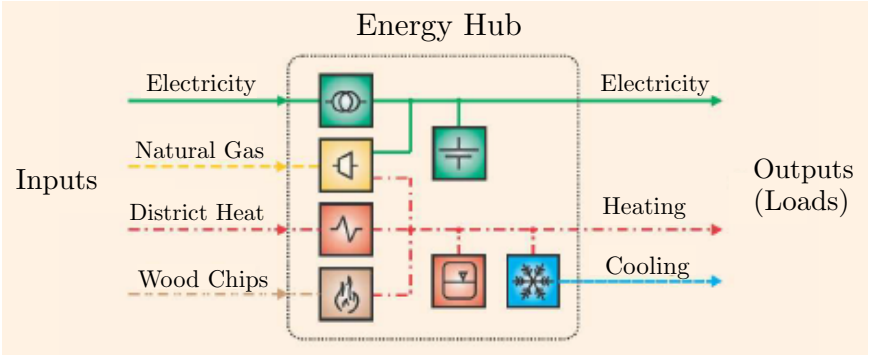

Figure 1.1: Sketch of energy hub in a utility network (Geidl and Andersson, 2007).

To ensure the matching of energy supply and demands at all times, power grids coordinate a diverse set of energy systems like generators, loads, storage devices, wind turbines, etc. at multiple timescales (from hours to milliseconds). Such coordination is achieved through hierarchical (multiscale) market transactions operated by the Independent System Operators (ISOs). Faster and more volatile energy markets provide many incentives for electricity generation and consumption systems (e.g., manufacturing plant, building). Large-scale battery systems and building systems are becoming key providers of dynamic flexibility to the power grid (Hao et al., 2012; Fares et al., 2014). As the level of coupling between gas, electric infrastructures, and renewable sources of power such as wind and solar is increasing, introducing new forms of dynamic behavior in the networks, the system operators realize that better infrastructure coordination is required to mitigate the effects of extreme events and uncertainties (ISO New England, 2014). Common examples of extreme events affecting the energy infrastructure include weather events such as the polar vortex experienced during the winter in the midwestern U.S. For instance, the polar vortex of 2014 in the midwest that resulted in economic losses of $175 million per hour for the power plant operators in states like California, Massachusetts, and Texas due to gas

shortages (Helman, 2014; Kemp, 2014).

Therefore, there is a great interest in designing automation architectures that handle uncertainties in the markets and demands, and exploit the fast timescale energy markets. These coupled energy networks have a multiscale dynamics because of the presence of systems operating at multiple different frequencies such as the long-term periodicity in demands and real-time prices. The challenge associated with this, however, is that the control algorithms for these systems would result in large-scale optimization problems which can become intractable for the current optimization solvers. In this dissertation, we propose control frameworks to handle uncertainties in markets and demands that can be applied to various energy storage systems such as batteries, central HVAC plants for large campuses, buildings, etc. We implement the control schemes in `Julia` programming language (Bezanson et al., 2012, 2017) and leverage the algebraic modeling capabilities of the package `JuMP` (Dunning et al., 2017).

1.2 Outline

The goal of this dissertation is to address the challenges discussed in Section 1.1 by developing a multiscale model predictive control architecture for energy systems that handles interactions and uncertainties at multiple timescales.

The structure of this thesis is described below:

- In the next section of this chapter, we present a background on energy networks, electricity market operations, the role of energy storage systems in electricity markets, and model predictive control. Then, we present a preliminary electricity market case study for electrochemical storage systems using a deterministic optimization modeling framework.

- In Chapter 2, we present a stochastic model predictive control (MPC) framework to simultaneously determine the battery participation commitments in ISO energy and

frequency regulation markets and mitigate demand charges for an attached modulated load by accounting for the uncertainties explicitly in the optimization models.

- In Chapter 3, we present a computational framework that integrates forecasting, uncertainty quantification, and model predictive control (MPC) to benchmark the performance of deterministic and stochastic MPC.

- In Chapter 4, we conduct detailed closed-loop simulations and systematic benchmarks for stochastic MPC framework for a central heating, ventilation, and air conditioning (HVAC) plant of a typical university campus using real data to forecast and quantify the uncertainty of disturbances affecting the system over multiple timescales (electrical loads, heating/cooling loads, and energy prices).

- In Chapter 5, we develop a hierarchical model predictive control (MPC) framework using stochastic programming to handle long horizon (or infinite horizon) problems for periodic systems by solving a retroactive optimization problem that progressively accumulates historical data to deliver the optimal periodic state targets.

- In Chapter 6, we provide a scalable approach to handle complex MPC applications with uncertainties evolving over long time horizons and with fine time resolutions by deriving and interpreting stochastic dual dynamic programming (SDDP) from the perspective of MPC.

- In Chapter 7, we propose a dual dynamic programming based approach for solving multiscale mixed-integer model predictive control (MPC) problems that can handle real constraints of energy systems such as discrete actuators and minimum startup loads.

- In Chapter 8, we present a summary of the major accomplishments of this dissertation. Some future directions arising from this dissertation are also discussed.

1.3 Background

1.3.1 *Energy Networks*

Energy network coupling introduces rich dynamic behavior. Electrical networks have very fast dynamics, while gas, water, and steam pipelines have slow dynamics because of transportation delays. Multiple timescales are introduced in the system because of these factors and also from disturbances (e.g. gas/electricity demands) that drive different networks (Baldea and Daoutidis, 2007; Jogwar et al., 2009; Zavala, 2014). The presence of multiple timescales complicates the coordination of energy networks (Shahidehpour et al., 2005). In Chiang and Zavala (2016), the authors highlight the benefits of coordinating gas and electric transmission networks with a case study in Illinois. Under coordinated control, gas targets of power plants could be tracked perfectly and 7% more gas was delivered to the power plants which amounts to $1,070,000 of additional gas in a single day. The increased gas delivered also resulted in increased revenue for the power plants by 27%. The increased flexibility was the result of enabling power plants to be operated as controllable gas demands.

The large-scale problems arising from these dynamical systems require reduced-order representations to make them tractable for real-time control applications. Such reduced representations can be developed by separation of timescales (Bıyık and Arcak, 2008). In Simon and Ando (1961), the authors established that there exist groups of strongly interacting subsystems within a network and the interaction between subsystems in the same group is much stronger than the interaction of subsystems from different groups. They also found relationships between weak network connections and timescales. Dynamic behavior of linear systems has also been analyzed over multiple timescales in a study (Coderch et al., 1983), in which they also provide an aggregation method to produce simplified models that are valid at progressively slower timescales. Sparsity and density of network interconnections also influence the separations of timescale (Bıyık and Arcak, 2008; Chow and

Kokotovic, 1985). The authors in Pappas et al. (2000) have created consistent aggregated characterizations for dynamical systems in which the controllability of the aggregated system is guaranteed if and only if the full system is controllable. The analysis tools resulting from these works have been applied to economic systems, Markov chains, biology, and power systems (Kokotovic, 1980; Delebecque et al., 1984; Chow and Kokotovic, 1985; Simon and Ando, 1961; Bıyık and Arcak, 2008).

1.3.2 *Electricity Market Operations*

Understanding the economic incentives provided by generation and load flexibility requires careful consideration of the wholesale electricity market structures and diverse products. The electricity markets are governed by ISOs (Independent System Operators) such as California ISO (CAISO), Pennsylvania-New Jersey-Maryland (PJM) Interconnection, and Midcontinent ISO. Markets are structured at multiple time levels, namely day-ahead (hourly market commitments) and real-time markets (commitments ranging from minutes to seconds). In day-ahead markets (DAM) the electricity is traded in intervals of 1 hour, real-time markets (RTM), on the other hand, can have varying timescales depending on the ISO operating it. Real-time market is a spot market in which utilities can buy power to meet the last few increments of demand not covered in their day-ahead schedules. The due times for the bids also vary from one ISO to the other. For example, in CAISO, the real-time markets are operated at two timescales, a 15-minute timescale which is commonly called as quarter-hourly market (QHM) or fifteen-minute market (FFM), and a 5-minute timescale which is called as real-time market (RTM), although both are categorized under the real-time market operations by CAISO. In CAISO, the DAM bids are due no later than 10 a.m. the day before while the QHM and RTM bids are due 75 minutes for the start of each trading hour. The typical time-varying energy prices for day-ahead and real-time markets in CAISO are shown in Figure 1.2. The RTM is more volatile, and at times the prices can even go negative here. This provides an opportunity to the fast dy-

namic energy systems to capture the negative spikes in prices and maximize their revenue potentials.

The power grid receives additional (contingency) flexibility via regulation and reserve ancillary services from the electricity markets. Energy systems providing regulation capacity permit the power Automatic Generator Control (AGC) layer run by the ISO or similar grid entity to adjust their load set-point within a specified band of power (Jaleeli et al., 1992) at a frequency of every 2 to 15 seconds. The regulation service provider is compensated both for the amount of regulation capacity (a load flexible band is offered) and the amount of mileage (performance of tracking the AGC signal). Spinning reserves provide an additional safeguard against unplanned outages and increased loads. Spinning reserves rarely need to be dispatched and resources providing reserves are compensated for providing flexibility/contingency. The requirements for ancillary services are expected to grow because of the high-frequency variations in the power grid coming from the intermittent and non-dispatchable wind and solar power (Walling et al., 2008). With the retirement of coal-fired generators, reductions in the supply of ancillary services are expected, creating additional opportunities for flexible load providers (Kirby et al., 2011). Figure 1.3 shows an illustration of the electricity market operations with various kinds of energy technologies participating in the electricity markets to provide various services to the power grid.

Market participation has been studied extensively for a variety of resources such as combined heat and power (CHP) plants (Rong and Lahdelma, 2007; Al-Mansour and Kožuh, 2007; De Paepe and Mertens, 2007; Christidis et al., 2012; Mitra et al., 2013), electrochemical manufacturing facilities (Babu and Ashok, 2008), HVAC systems for large buildings (Hao et al., 2012; Zhao et al., 2013, 2015; Baccino et al., 2014), and manufacturing systems with thermal energy storage (Feng et al., 2015). Many studies have also analyzed the economics of energy storage systems (ESS) such as flywheels (Walawalkar et al., 2007; Bradbury et al., 2014), batteries (Oudalov et al., 2007a; Mercier et al., 2009; Ekman and Jensen, 2010; Walawalkar et al., 2007; Sortomme and El-Sharkawi, 2012; Fares et al., 2014; Khalilpour and Vassallo, 2016; Bradbury et al., 2014; Zakeri and Syri, 2014), and concen-

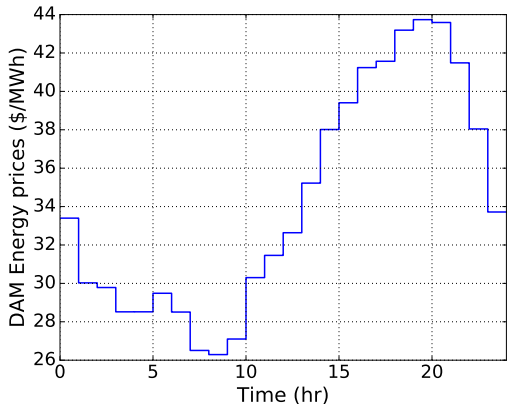

(a) Day-ahead market

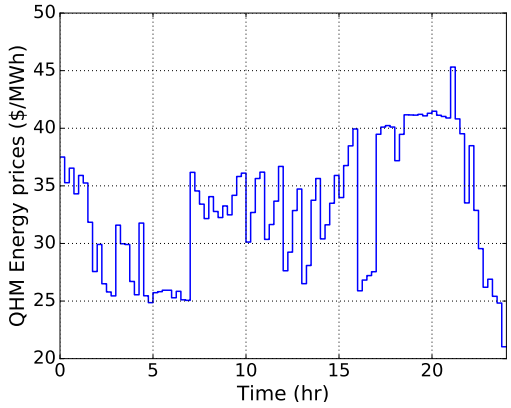

(b) Quarter-hourly market

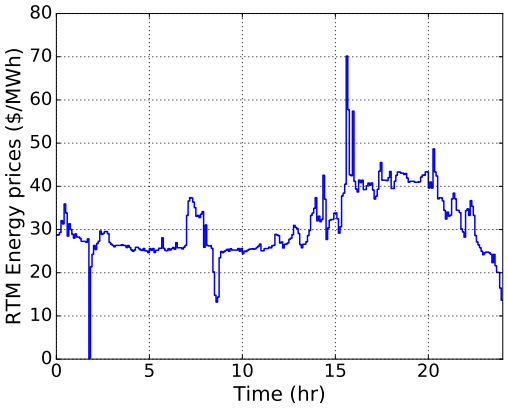

(c) Real-time market

Figure 1.2: Energy prices for a day in the markets in California. The day-ahead market (DAM) prices vary at 1-hour timescale, the quarter-hourly market (QHM) prices are at 15-minute timescale, the real-time market (RTM) prices are at 5-minute resolution.

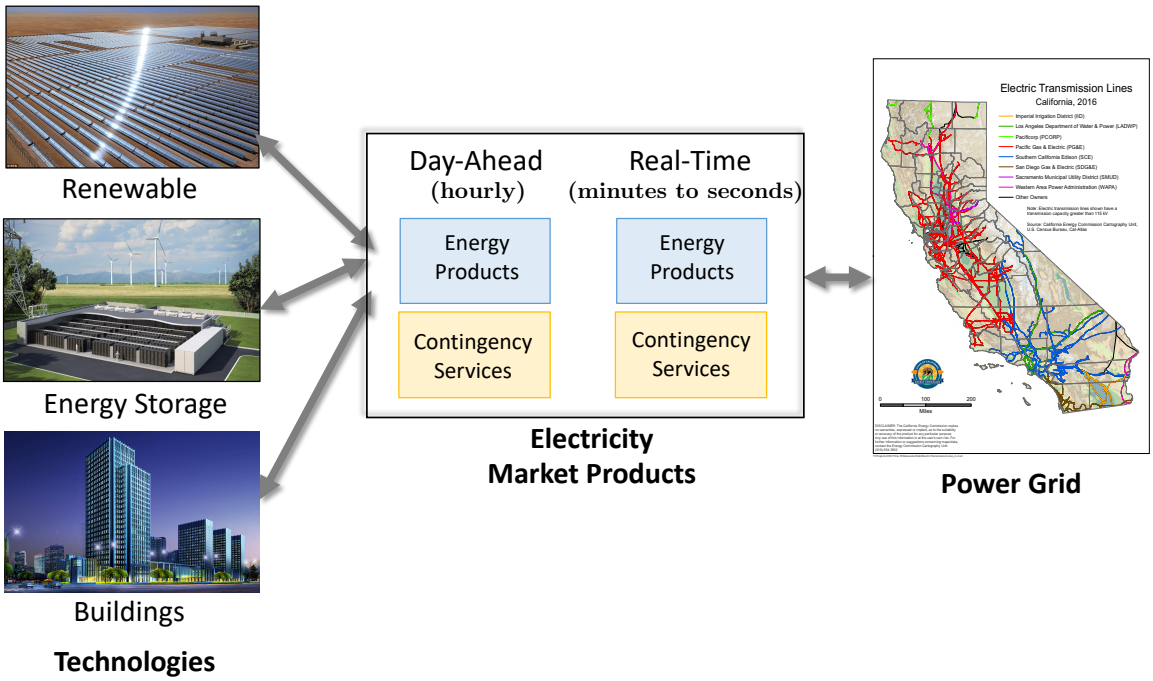

Figure 1.3: Illustration of the electricity market interactions with various energy technologies.

trated solar power generators with thermal energy storage (Kost et al., 2013; Madaeni et al., 2011; Lizarraga-Garcia et al., 2013; Usaola, 2012) interacting in wholesale electricity markets. Several studies focus on only energy transactions in day-ahead markets (Sioshansi et al., 2009; Kost et al., 2013; González et al., 2014; Usaola, 2012; Brunetto and Tina, 2007; Dicorato et al., 2012; Bradbury et al., 2014), while some other papers consider simultaneous sale of energy and ancillary services in the day-ahead market alone (Walawalkar et al., 2007; Sortomme and El-Sharkawi, 2012; Ekman and Jensen, 2010; Madaeni et al., 2011; He et al., 2015) or analyze only regulation revenues (Fares et al., 2014; Oudalov et al., 2007a). The economics of participation in the real-time markets are considered in only a few studies (Sarker et al., 2015; Bradbury et al., 2014; Zakeri and Syri, 2014; Lizarraga-Garcia et al., 2013; Chen and Garcia, 2016). In Sarker et al. (2015), the authors study bids for energy and reserves from electric vehicle aggregators in day-ahead markets while real-time prices are only considered to settle reserve dispatches. In Zakeri and Syri (2014), the authors compare revenue opportunities for five energy storage technologies, and find 37-300% higher

revenues in the real-time markets.

1.3.3 *Energy Storage in Electricity Markets*

Battery energy storage systems (BESS) are increasingly becoming important for the applications such as providing frequency regulation to the electric grid and for peak shaving applications (Rastler, 2010; Oudalov et al., 2006; Donadee and Ilić, 2014). Compared with traditional generators, the battery storage systems have much higher ramping capabilities and they can respond faster with better performance (He et al., 2015). Economic benefits of utilizing BESS for frequency regulation have been studied previously using heuristic unit commitment algorithms (Alt et al., 1997) and using optimization methods (de Salis et al., 2014; He et al., 2015; Sigrist et al., 2013; Lucas and Chondrogiannis, 2016). Several papers have also studied the economic potential of the energy storage in the regulation markets by using hybrid electric vehicles or vehicle-to-grid technologies (Foster and Caramanis, 2010; Han et al., 2010; Lin et al., 2014; White and Zhang, 2011). Stochastic optimization methods, that capture uncertainty in the market price and regulation signals, have also been used to maximize the benefits of using energy storage systems for frequency regulation (Donadee and Ilić, 2012; Donadee and Ilić, 2014; Vagropoulos and Bakirtzis, 2013). Bulk energy storage systems have also been identified for peak shaving among large industrial plants and power utilities (Oudalov et al., 2007b). The electricity bill for any facility consists of payment for energy and for the peak power demand. The payment for the peak power demand, also known as the demand charge, corresponds to the highest power demand during a specific time period (typically a month) and often contributes as much as half of the plant's electricity bill (Oudalov et al., 2007b). Thus, they would like to minimize their peak demand over that period using on-site generation or storage. This is also known as "peak shaving" (Oudalov et al., 2007b). A battery can utilize its storage capability to effectively reduce the energy demand of a plant or a facility from the grid during peak load durations (see Figure 1.4). This can help the plant to reduce the peak demand and

thus reduce the electricity bill. From previous research and installations (Alt et al., 1997; de Salis et al., 2014; Johnson et al., 2011; Leadbetter and Swan, 2012; Lucas and Chondrogiannis, 2016; Oudalov et al., 2006, 2007b; Rahimi et al., 2013; Sebastián, 2016), it is evident that battery storage systems can be used for peak load shaving application effectively and profitably.

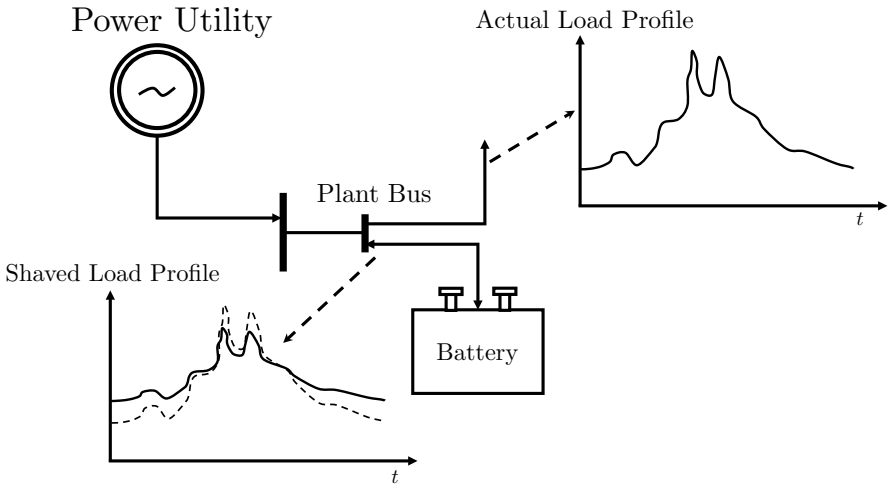

Figure 1.4: Peak shaving application of energy storage systems (Figure adapted from Oudalov et al. (2007b)).

A few studies have addressed simultaneous frequency regulation and peak shaving applications of energy storage systems (Alt et al., 1997; Sigrist et al., 2013; White and Zhang, 2011; Lucas and Chondrogiannis, 2016). However, they consider a deterministic setting to study the economic benefits of batteries in these applications which may not be appropriate to plan a battery usage policy for future (e.g., bidding in markets). To address these issues, we propose a stochastic model predictive control framework in Chapter 2 to study the economic benefits of the battery by planning a usage policy under uncertainty in the markets and building loads.

1.3.4 *Model Predictive Control*

Model predictive control (MPC) or receding horizon control, is an optimization-based control technology that relies on real-time solutions to optimal control problems. The optimal control problem contains the dynamics and physical constraints of the system. In such an optimization setting it is also possible to handle complex nonlinear dynamics and imposition of explicit constraints on the system (Rawlings et al., 2017; Camacho and Alba, 2013). There are also predictive models in the controller formulation to enable anticipation of future events like disturbances. The controller deals with the prediction errors from the model and disturbance by recursively shifting the horizon of the controller at each time step. The control action obtained at each time step is fed to the system and the state of the system is updated before moving to the next time step, thus creating a feedback. MPC has been implemented in a variety of complex industrial systems such as chemical manufacturing, aerospace, buildings, and power plants (Qin and Badgwell, 2003, 2000; Ma et al., 2012b; Oldewurtel et al., 2012; Lopez-Negrete et al., 2013). The inherent system dimensionality controls the computational complexity of MPC. The presence of nonlinear dynamics, and in many cases, the need to expand prediction horizons and capture fast and slow dynamics also leads to greater computational complexity (Baldea and Daoutidis, 2007). The presence of multiple timescales is particularly relevant because this dictates the length of the prediction horizon and the update frequency (Scattolini, 2009). Expanding the prediction horizon is also necessary to manage short- and long-term storage in energy systems (Zavala et al., 2009b,a; Lefort et al., 2013).

An extension of MPC to deal with uncertainty in the variables is known as stochastic model predictive control (stochastic MPC) (Cannon et al., 2007; Kouvaritakis and Cannon, 2013). The stochastic MPC takes explicit account of the probability distribution of the stochastic model uncertainty in the optimization of predicted performance, and at each time step of the MPC, a stochastic optimization problem needs to be solved. Stochastic MPC has recently been implemented in several applications like for semi-autonomous ve-

hicles with an uncertain driver model (Gray et al., 2013), microgrid management (Hooshmand et al., 2012), building climate control (Oldewurtel et al., 2013).

1.4 An Electricity Market Case Study

In this section, we present a preliminary case study using optimization models for the participation of battery energy storage systems (BESS) in the multiscale energy markets in a deterministic setting. Using a multiscale optimization model, we develop a framework for the economic analysis for various applications of BESS such as energy arbitrage and frequency regulation.

1.4.1 *Modeling and Economic Analysis of Participation in a multiscale Energy Market in Deterministic Setting*

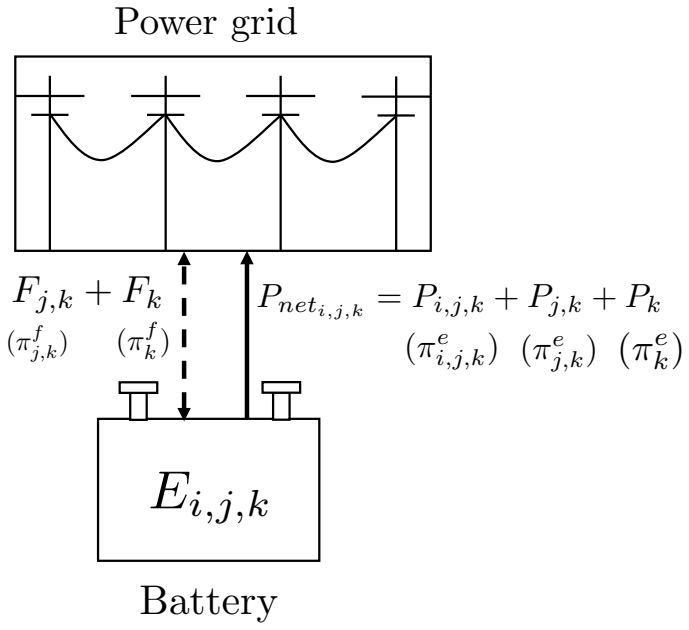

Figure 1.5: Illustration of transactions of energy between the power grid and battery system. Quantities in parentheses denote the price for the corresponding transaction.

We formulate an optimization problem to maximize revenue for the battery system

participating in the California electricity markets in a deterministic setting, assuming no uncertainty in any of the market and demand signals (Dowling et al., 2017). In this formulation, the battery system participates at three timescales of the CAISO markets, the day-ahead market (DAM), the fifteen-minute market or quarter-hourly market (FFM or QHM), and the 5-minute real-time market (RTM).

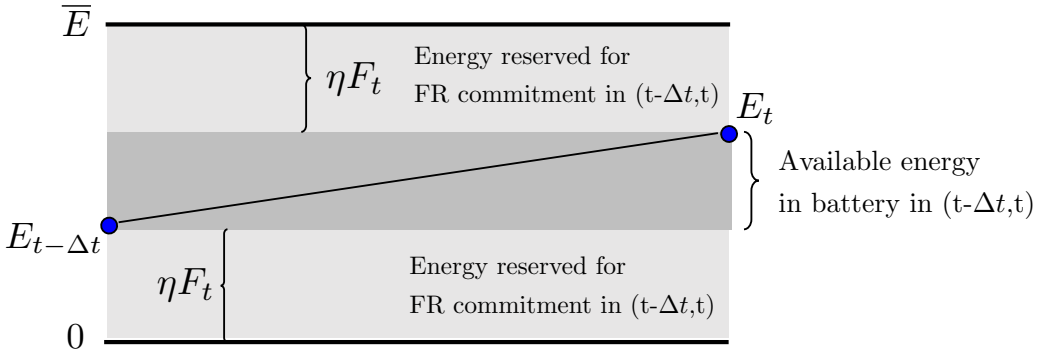

Figure 1.6: Illustration of the reserved energy storage for regulation capacity (given by constraints (1.4.5) and (1.4.6)).

In the optimization model for the system, we define the sets of time indices $\mathcal{T}_\mathcal{D} := \{1, .., n_\mathrm{d}\}$ corresponding to the hourly intervals of day-ahead market with (n_d can be the number of hours of planning period, say 24 for 1-day planning), $\mathcal{T}_\mathcal{Q} := \{1, .., n_\mathrm{q}\}$ corresponding to the time intervals in quarter-hourly market within each day-ahead hour (so, $n_q = 4$) and $\mathcal{T}_\mathcal{R} := \{1, .., n_\mathrm{r}\}$ corresponding to the intervals in real-time market within each QHM interval (so, $n_r = 3$).

A time instance for variables in different timescales of the market is denoted as follows: with a lone subscript index $k \in \mathcal{T}_\mathcal{D}$ for the DAM timescale corresponding to the 1-hour interval between hours $k-1$ and k; $(j,k) \in \mathcal{T}_\mathcal{Q} \times \mathcal{T}_\mathcal{D}$ for the QHM timescale corresponding to the j^{th} 15-minute interval between hours $k-1$ and k; and $(i,j,k) \in \mathcal{T}_\mathcal{R} \times \mathcal{T}_\mathcal{Q} \times \mathcal{T}_\mathcal{D}$ for the RTM timescale corresponding to the i^{th} 5-minute interval within the j^{th} 15-minute interval of the k^{th} hour. For example, the state of charge (energy level) of the battery in real time will be denoted by $E_{i,j,k}$, while the power sold by the battery to the day-ahead market will be denoted by P_k. The transactions of energy have been illustrated with a schematic diagram

in Figure 1.5.

The physical and market operational constraints are described for the battery in Equations (1.4.1)-(1.4.8). In the model, η is the efficiency of the battery (0.95 for this case study) and Δt is the duration of a real-time interval (i.e. 5 min). Net discharge (net power sold, P_k in DAM, $P_{j,k}$ in QHM and $P_{i,j,k}$ in RTM) from the battery at any time has to be less than the maximum discharge capacity of the battery, $\overline{P}$ (1 MW in this case), and is given in Equation (1.4.1). The frequency regulation capacity F_k is sold in DAM and $F_{j,k}$ is sold in QHM (regulation market does not operate at the 5-minute scale in CAISO). The regulation capacities will limit the net discharge capacity of the battery, given by constraint (1.4.2) and (1.4.3). The dynamics of the state of charge (energy level) of the battery ($E_{i,j,k}$) is given by Equation (1.4.4), where $E_{0,1,1} = \overline{E}$ (the storage capacity which is also the initial storage level, 1 MWh in this study), $E_{0,j,k} = E_{n_{r,j-1,k}}$ for $(j,k) \in \mathcal{T}_\mathcal{Q} \setminus \{1\} \times \mathcal{T}_\mathcal{R}$. The constraints (1.4.5) and (1.4.6) are used to ensure sufficient storage capacity to accommodate regulation sales (illustrated in Figure 1.6). We also put a bound on the ramping of discharge rate of the battery by constraints (1.4.7), but these bounds are set to large values to reflect fast dynamics of battery. Bounds on the decision variables are given in Equations (1.4.8) and finally, the objective of optimization, Φ is given by the total revenue from the market participation which we want to maximize (Equation (1.4.9)). All constraints (1.4.1)-(1.4.8) are written for all $i \in \mathcal{T}_\mathcal{R}, j \in \mathcal{T}_\mathcal{Q}$ and $k \in \mathcal{T}_\mathcal{D}$.

$$-\overline{P} \leq P_{net_{i,j,k}} = P_k + P_{j,k} + P_{i,j,k} \leq \overline{P} \tag{1.4.1}$$

$$P_{net_{i,j,k}} + F_k + F_{j,k} \leq \overline{P} \tag{1.4.2}$$

$$P_{net_{i,j,k}} - F_k - F_{j,k} \leq \overline{P} \tag{1.4.3}$$

$$0 \leq E_{i,j,k} = E_{i-1,j,k} - \frac{1}{\eta} P_{net_{i,j,k}} \Delta t \leq \overline{E} \tag{1.4.4}$$

$$E_{i,j,k} + \eta (F_k + F_{j,k}) \Delta t \leq \overline{E} \tag{1.4.5}$$

$$E_{i,j,k} - \frac{1}{\eta} (F_k + F_{j,k}) \Delta t \geq 0 \tag{1.4.6}$$

$$-\overline{\Delta P} \leq P_{net_{i,j,k}} - P_{net_{i-1,j,k}} \leq \overline{\Delta P} \tag{1.4.7}$$

$$-\overline{P} \leq P_{i,j,k}, P_{j,k}, P_k \leq \overline{P}, \quad 0 \leq F_{j,k}, F_k \leq \overline{P} \tag{1.4.8}$$

$$\Phi = \sum_{i \in \mathcal{T}_{\mathcal{R}}} \sum_{j \in \mathcal{T}_{\mathcal{Q}}} \sum_{k \in \mathcal{T}_{\mathcal{D}}} (\pi^e_{i,j,k} P_{i,j,k} + \pi^e_{j,k} P_{j,k} + \pi^e_k P_k + \pi^f_{j,k} F_{j,k} + \pi^f_k F_k) \Delta t \tag{1.4.9}$$

1.4.2 *Results and Economic Analysis in California ISO (CAISO)*

We performed a case study that considers the participation of a simple battery storage system in CAISO markets. Table 1.1 presents the revenue estimates for the battery system under different operating policy restrictions. Performing only energy arbitrage in all three market layers yields 139.1 k$/year in net revenue. Adding ancillary services participation increases net revenues to 199.6 k$/year (44% increase). Interestingly, restricting participation to only the day-ahead market reduces net revenues substantially: 10.5 k$/year (energy only) and 72.8 k$/year (energy and ancillary services). Thus with only DAM participation, also providing ancillary services increases revenues by 600% relative to only energy arbitrage. These results also highlight the importance of studying revenues available at multiple market layers, especially at fast timescales. In particular, most previous studies for batteries only estimate revenues using day-ahead market prices (Fares et al., 2014; Walawalkar et al., 2007; Sortomme and El-Sharkawi, 2012; Ekman and Jensen, 2010; Dicorato et al., 2012; Sarker et al., 2015), and thus miss important economic opportunities.

Table 1.1: Revenues for different market participation schemes and operating mode combinations. Here, RTM includes both 15-minute and 5-minute timescales.

	DAM + RTM	**DAM only**	**RTM only**
Energy only	139.1 k$/year 100%	10.5 k$/year 8 %	115.0 k$/year 83%
Energy & Regulation	199.6 k$/year 100%	72.8 k$/year 36 %	141.8 k$/year 71%

We also examined the total energy transactions (sales and purchases) under the six market participation schemes. Interestingly, with participation in multiple timescales, the dominant trend is to purchase energy in the RTM layer and resell it in the slower layers (DAM and QHM), because typically, the average prices in markets at faster timescales are

lower than slower timescale. Because battery systems have much faster dynamics than conventional generators, they are capable of exploiting opportunities in the RTM (fastest market layer). This study thus highlights revenue opportunities for electricity storage systems from markets at faster timescales. Table 1.1 also shows that the revenues from simultaneous ancillary services and energy participation are 43% higher than those obtained with energy participation only (revenue increases from 139.1$/year to 199.6$/year). This, again, illustrates how the proposed framework can be used to identify which market layers and products offer the greatest economic potential.

1.4.3 *Summary*

We presented a case study of battery energy storage systems in the electricity markets using a deterministic optimization modeling framework. Based on the case study, we conclude that the energy storage systems can bring significant revenues by participating in the electricity markets and providing frequency regulation services to the power grid. It is important to note, however, that this case study did not account for the significant uncertainties in electricity prices and loads, that cause uncertainties to revenue streams that can be generated from different markets. In Chapter 2, we will present a stochastic model predictive control framework that handles uncertainty explicitly in the optimization model and determines the battery participation commitments in the ISO energy and frequency regulation markets and simultaneously mitigate demand charges for an attached modulated load.

2

STOCHASTIC MODEL PREDICTIVE CONTROL FOR STATIONARY BATTERY SYSTEMS

In this chapter, we propose a stochastic model predictive control (MPC) framework to handle uncertainties in the electricity markets and demands. The proposed stochastic MPC framework determines the battery participation commitments in ISO energy and frequency regulation markets, and simultaneously mitigates demand charges for an attached modulated load (Kumar et al., 2018a). The proposed framework solves a two-stage stochastic program that maximizes the expected revenue over a receding horizon and that factors in uncertainty of the load, energy and regulation prices, and regulation signals. We propose to use a Ledoit-Wolf covariance estimator to generate load and price scenario profiles from limited historical data. We benchmark the performance of stochastic MPC against perfect information MPC and deterministic MPC for different prediction horizon lengths and demand charge discounting strategies by using real load data for a typical university campus and price and FR signal data from PJM.

2.1 Introduction

Battery systems are flexible assets that can provide energy and frequency regulation (FR) capacity for independent system operators (ISOs) and aid utility companies by providing demand-side management capabilities for energy-intensive facilities (e.g. buildings or manufacturing) (Rastler, 2010; Oudalov et al., 2006). This flexibility is becoming increasingly valuable as more intermittent renewable power is injected into the grid (Donadee and Ilić, 2012; Donadee and Ilić, 2014). In particular, ISOs have reported an increased demand for fast dispatchable resources and frequency regulation services to mitigate high-frequency renewable fluctuations and stranded power (Kim et al., 2016). Battery systems can also be strategically placed in the network to enhance system-wide performance (He et al., 2015; Sioshansi et al., 2009).

The economic benefits of using stationary battery systems to provide FR have been studied in Mercier et al. (2009); He et al. (2015); Mohsenian-Rad (2016); Kottick et al. (1993); Singh et al. (2014). The focus of Mercier et al. (2009); He et al. (2015); Mohsenian-Rad (2016) is on exploring the economic opportunities of using battery systems to provide FR capacity in day-ahead markets. The studies in Kottick et al. (1993) and Singh et al. (2014) propose real-time control strategies for providing FR services to the grid. The potential of harnessing the flexibility of batteries from collections of electric vehicles has also been studied in Foster and Caramanis (2010); Han et al. (2010); Lin et al. (2014); White and Zhang (2011). These studies use optimization methods to obtain bidding strategies in day-ahead regulation markets. The economic benefits of using battery systems to reduce demand charges in buildings and microgrids (via peak shaving) have been studied in de Salis et al. (2014); Johnson et al. (2011); Leadbetter and Swan (2012); Oudalov et al. (2007a); Rahimi et al. (2013); Lu et al. (2014); Dong et al. (2011). Some of these studies also explore battery sizing issues (Oudalov et al., 2007a; Leadbetter and Swan, 2012; Lu et al., 2014; Mercier et al., 2009). A common limitation of these studies is that they assume that the battery either provides FR or demand-side management capabilities.

The use of batteries for *simultaneous* FR and demand charge mitigation has been studied in Alt et al. (1997); Sigrist et al. (2013); White and Zhang (2011); Lucas and Chondrogiannis (2016); Sebastián (2016); Oudalov et al. (2006). In such settings, the battery has the dual goal of collecting revenue from FR markets while modulating the load of an attached load to mitigate demand charges. In particular, the authors in Sigrist et al. (2013) and Oudalov et al. (2006) provide comparative studies to show that the *value of stationary batteries* can be improved significantly by providing both services. These studies use optimization models that assume perfect knowledge of all the market and load parameters. Consequently, while these approaches are valuable for design and planning tasks, they can be limited in real-time operational settings where diverse uncertainties can hinder performance.

Stochastic optimization models have been recently proposed to maximize the economic performance of batteries from electric vehicles while capturing uncertainty in market prices and FR signals (Donadee and Ilić, 2012; Donadee and Ilić, 2014; Vagropoulos and Bakirtzis, 2013; Moura et al., 2010; Foster and Caramanis, 2010; Wang et al., 2016). Some of these studies use stochastic dynamic programming (stochastic DP) methods to optimize charging and FR commitments (Foster and Caramanis, 2010; Donadee and Ilić, 2012; Donadee and Ilić, 2014; Moura et al., 2010). In Donadee and Ilić (2012) and Donadee and Ilić (2014), the authors model uncertainties in prices and FR signals as Markov random variables (this is needed to apply stochastic DP methods) and demonstrate that their proposed method results in lower charging costs than a deterministic model predictive control (MPC) strategy. In Vagropoulos and Bakirtzis (2013), optimal bidding strategies for an electric vehicle aggregator participating in day-ahead energy and FR markets are determined using stochastic optimization. Here, uncertainties are captured by deriving realizations from past historical data. All of these studies consider only mobile battery systems. Demand-side management in a commercial building using both mobile and stationary batteries has been studied in Wang et al. (2016). Here, the authors show that a stochastic optimization approach results in higher cost savings than its deterministic counterpart. This study is limited in that it conducts these studies in the context of day-ahead

planning and for one day of operation. In a real setting, however, the peak demand is charged by the utility company on a monthly basis and one must capture real-time behavior.

In summary, to the best of our knowledge, stochastic optimization techniques have not been used in stationary battery systems that simultaneously participate in real-time energy and FR markets and perform demand-side management. Moreover, existing studies provide limited comparisons with deterministic approaches and have not explored important factors affecting performance such as prediction horizon lengths and long-term peak demand charges. Such information can be valuable to industrial control vendors, which are seeking to reduce costs and complexity of battery control technologies. Moreover, as in general stochastic optimization applications, uncertainty modeling, and scenario generation from historical and operational data remains a challenge. This is particularly important in batteries that provide dual functions because prices and loads tend to exhibit multiscale correlations.

In this work, we propose a stochastic MPC framework to determine optimal participation strategies in energy and FR markets for stationary batteries while mitigating demand charges of an attached load. We propose to use the Ledoit-Wolf covariance estimator to generate realizations that capture multiscale correlations from historical data. We use the framework to study the flexibility and economic benefits provided by a one-MW battery system attached to an aggregated load from a collection of buildings. We use the proposed framework to study the benefits of stochastic MPC policies over those obtained with standard deterministic MPC and perfect information strategies. We study the effect of the prediction horizon length and demand charge discounting on the performance of the MPC policies. Using real load data for a typical university campus, we find that stochastic MPC can recover 83% of the ideal value of the battery (obtained under perfect information) while deterministic MPC can only recover 73%.

2.2 Problem Formulation

We begin by describing the decision-making setting under which the battery is operated. We also describe a deterministic MPC model that will serve as a basis for comparison.

2.2.1 *Decision-Making Setting*

The goal is to determine the optimal participation strategies for the battery in energy and FR markets operated by ISO (here we consider PJM) while simultaneously mitigating demand charges from a utility company. We describe this setting on a real-time planning level where the energy and regulation commitments are considered on hourly intervals. The various cost and revenue components that are considered in the optimization formulation are:

- *Energy Transactions (hourly)*: The battery system purchases energy from the grid to recharge and discharges to provide energy for the building load and for the regulation signal. The energy transactions are charged at the real-time energy price.

- *Frequency Regulation Capacity (hourly)*: The ISO compensates the battery system for providing an operational band (compensated based on regulation capacity prices) around a charge/discharge power level (charged at real-time energy prices). The ISO can request the battery to dispatch a fraction of the committed regulation capacity based on the grid requirements (e.g., at every 2 seconds in PJM). The real-time FR dispatch signal from the ISO is a zero-mean signal with a bounded range of [-1,+1]. In our planning setting, we consider an hourly average FR signal. The hourly-averaged dispatch signal is also a zero-mean signal with a significantly smaller variance. In particular, the variance of the FR signal is reduced due to averaging. A typical 2-second dispatch signal from PJM along with its hourly average is shown in Figure 2.1.

- *Demand Charges (monthly)*: The attached load is charged for the peak demand (at a fixed demand charge price) incurred over a month by the utility company.

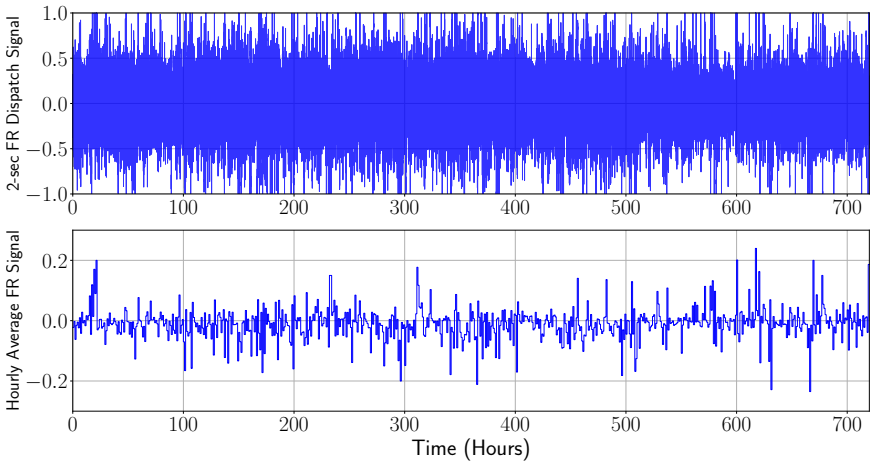

Figure 2.1: FR signal (real-time and hourly average) from PJM.

According to the rules of the energy and FR markets set by the ISO, the decision of allocating commitments for a trading hour can be changed 90 minutes before the start of that trading hour (at the latest). Consequently, the battery has the opportunity to reschedule commitments for the subsequent trading hours after the first hour over which we commit. Because of this flexibility, we can solve an optimization problem every hour over a receding horizon to maximize net revenue based on price and load conditions observed in real-time. In the proposed receding horizon scheme, a battery charge/discharge policy is obtained by solving an optimization problem after every hour, starting 90 minutes before the start of the trading hour (i.e., the loads and prices need to be forecasted 90 minutes in advance). For simplicity in the presentation, we assume that *forecasts are obtained one hour in advance*, and thus systems conditions are uncertain over the next immediate hour and at subsequent times.

2.2.2 *Deterministic MPC Framework*

The combined battery and modulated load system is illustrated in Figure 2.2. This schematic illustrates the interaction between the system, the ISO, and the utility company. The de-

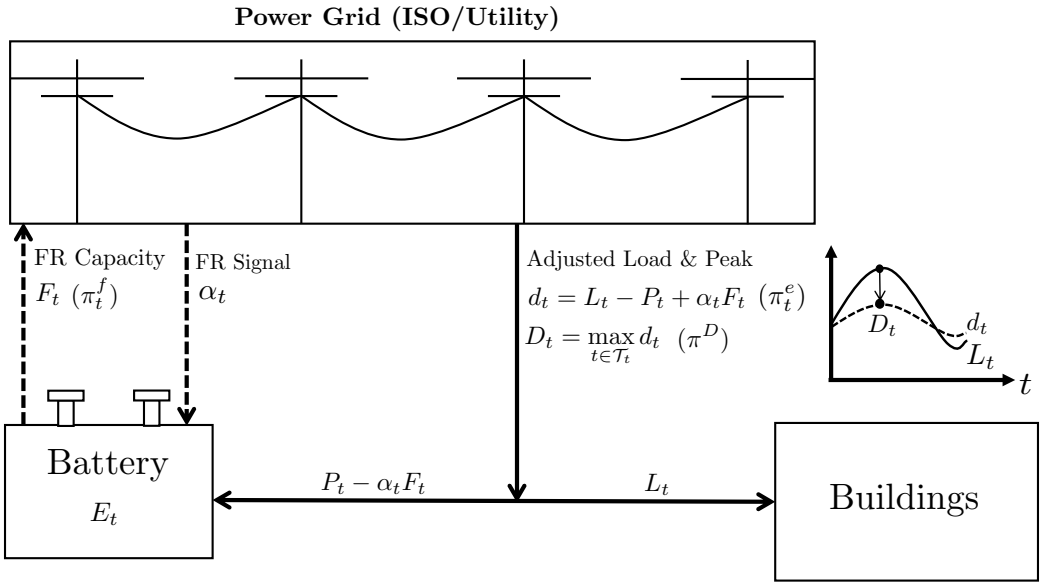

Figure 2.2: Interactions between battery system, ISO, and utility.

terministic MPC scheme solves a discrete-time optimal control problem at real-time t over the prediction horizon $\mathcal{N}_t := \{t+1, t+2, \ldots, t+N\}$ to maximize net revenue (minimize cost). Here, N is the length of the prediction horizon and index $t \in \mathbb{Z}_+$ corresponds to the time instant $t \cdot h$, where $h = 1$ hour. The optimization problem at time t uses forecasts for prices and loads over the prediction horizon $\mathcal{N}_t$. This scheme is run over a period that covers a month $\mathcal{M} := \{0, \ldots, M\}$, with $M = 720$.

We now proceed to describe the elements of the optimization problem solved in deterministic MPC. We then extend this to a stochastic formulation that factors in uncertainty. The following are the model parameters, data, and variables.

Model Parameters and Data

- $L_t \in \mathbb{R}_+$: Attached load [kW].
- $\hat{L}_t \in \mathbb{R}_+$: Forecast for load [kW].
- $\pi_t^e \in \mathbb{R}$: Electricity price [\$/kWh].
- $\hat{\pi}_t^e \in \mathbb{R}$: Forecast of electricity price [\$/kWh].

- $\pi_t^f \in \mathbb{R}_+$: FR capacity price [\$/kW].
- $\hat{\pi}_t^f \in \mathbb{R}_+$: Forecast for FR capacity price [\$/kW].
- $\pi^D \in \mathbb{R}_+$: Demand charge [\$/kW].
- $\alpha_t \in [-1, 1]$: Hourly average fraction of FR capacity requested by ISO [-]. If $\alpha_t > 0$, the ISO sends power to the battery while if $\alpha_t < 0$ the ISO withdraws power.
- $\hat{\alpha}_t \in [-1, 1]$: Forecast for the hourly average fraction of FR capacity requested by ISO [-].
- $\overline{E} \in \mathbb{R}_+$: Battery capacity [kWh].
- $\overline{P} \in \mathbb{R}_+$: Maximum discharging rate [kW].
- $\underline{P} \in \mathbb{R}_+$: Maximum charging rate [kW].
- $\rho \in \mathbb{R}_+$: Minimum fraction of battery capacity reserved as a buffer for a unit FR capacity [kWh/kW].
- $\overline{\Delta P} \in \mathbb{R}_+$: Maximum ramping limit [kW/h].
- $\widehat{D}_t \in \mathbb{R}_+$: Peak load observed up to time $t \in \mathcal{M}$ [kW].

Model Variables

- $P_t \in \mathbb{R}$: Net battery charge/discharge rate [kW]. If $P_t > 0$, the battery is being discharged and if $P_t < 0$ the battery is being charged.
- $F_t \in \mathbb{R}_+$: FR capacity provided to ISO [kW].
- $E_t \in \mathbb{R}_+$: Battery state of charge (SOC) [kWh].
- $d_t \in \mathbb{R}_+$: Load requested from utility [kW].
- $D_t = \max_{k \in \mathcal{N}_t} d_k$: Peak load over horizon $\mathcal{N}_t$ [kW]

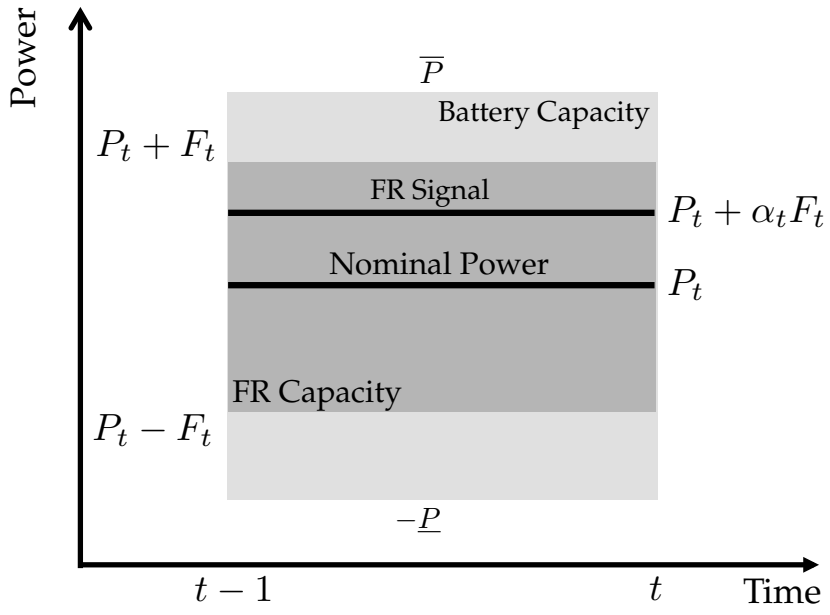

Figure 2.3: Sketch of FR capacity and FR dispatch signal.

In the proposed notation, all quantities with subindex t are *held constant* over the time interval $[(t-1)\cdot h, t\cdot h]$. The FR capacity F_t represents a *symmetric band* around the middle point P_t (the net charge/discharge rate). This is illustrated in Figure 2.3. In other words, the battery commits to a discharge level P_t and an FR band of total length $2\cdot F_t$. The actual power requested by the ISO for regulation (averaged over an hour) is $\alpha_t F_t$. A value of $\alpha_t = 1$ indicates that the ISO requests all the FR capacity by charging the battery while $\alpha_t = -1$ indicates that the ISO requests all the capacity by discharging the battery. When convenient, we use the compact notation $\pi_t = (\pi_t^e, \pi_t^f)$ and use $\pi_{\mathcal{N}_t}$ to denote the price trajectory over the horizon $\mathcal{N}_t$. We use a similar notation to denote trajectories for other quantities such as the price forecasts $\hat{\pi}_{\mathcal{N}_t}$, loads $L_{\mathcal{N}_t}$, load forecasts $\hat{L}_{\mathcal{N}_t}$, hourly average FR dispatch fractions $\alpha_{\mathcal{N}_t}$, and its forecasts $\hat{\alpha}_{\mathcal{N}_t}$. Based on this notation we highlight that the *uncertain quantities* (π_{t+1}^e, π_{t+1}^f, L_{t+1}, α_{t+1}) can only be forecasted an hour in advance (i.e., at time t).

Objective Function

In deterministic MPC we maximize the total *forecasted* net revenue (minimize the negative total cost) of the system over the prediction horizon $\mathcal{N}_t$:

$$\sum_{k\in\mathcal{N}_t} \hat{\pi}_k^f F_k - \sum_{k\in\mathcal{N}_t} \hat{\pi}_k^e d_k - \frac{\pi^D}{\sigma_t} \max_{k\in\mathcal{N}_t} d_k, \tag{2.2.1}$$

where $d_k = \hat{L}_k - P_k + \hat{\alpha}_k F_k$ is the *net residual demand*. The first term is the revenue obtained from the provision of FR *capacity*, the second term is the cost associated to purchasing power from the ISO, and the third term is the cost associated to the demand charge. The parameter σ_t is a *discounting factor* for the monthly demand charge price π_D, which can be used to adjust the demand charge when a prediction horizon of less than a month is used ($N < M$). As we will see, when such a discounting factor is not used, the MPC scheme can operate very conservatively because it will try to prevent the peak demand charge over the next immediate horizon $\mathcal{N}_t$.

We note that the terms $\hat{\pi}_k^e \hat{L}_k$ in the total net revenue (2.2.1) are constant quantities and thus do not affect the solution of the optimization problem. Consequently, we find it convenient to simplify the objective function as:

$$\sum_{k\in\mathcal{N}_t} \hat{\pi}_k^e (P_k - \hat{\alpha}_k F_k) + \sum_{k\in\mathcal{N}_t} \hat{\pi}_k^f F_k - \frac{\pi^D}{\sigma_t} \max_{k\in\mathcal{N}_t} d_k. \tag{2.2.2}$$

Constraints

The net charge/discharged battery power plus the FR capacity provided must be within the maximum discharging and charging rates $\overline{P}$ and $\underline{P}$:

$$P_k + F_k \leq \overline{P},\ k \in \mathcal{N}_t \tag{2.2.3a}$$

$$P_k - F_k \geq -\underline{P},\ k \in \mathcal{N}_t \tag{2.2.3b}$$

We note that these constraints make allowable battery charge/discharge limits a function of the committed FR capacity. In other words, the larger the FR commitment, the less capacity available to charge/discharge the battery.

The storage dynamics are given by the difference equation:

$$E_k = E_{k-1} - P_k + \hat{\alpha}_k F_k,\ k \in \mathcal{N}_t \tag{2.2.4}$$

The following constraint is used to ensure that a certain amount of energy is reserved for the committed FR capacity over the interval $(k-1, k)$:

$$\rho F_k \leq E_{k-1} \leq \overline{E} - \rho F_k,\ k \in \mathcal{N}_t \tag{2.2.5a}$$

$$\rho F_k \leq E_k \leq \overline{E} - \rho F_k,\ k \in \mathcal{N}_t \tag{2.2.5b}$$

These constraints impose a safety margin (cushion) to account for the fact that we do not capture the variability of the 2 second FR signal. Capturing variability at such high time resolutions would significantly increase the computational complexity of the optimization problem.

The battery ramp discharge rate is constrained as:

$$-\overline{\Delta P} \leq P_k - P_{k-1} \leq \overline{\Delta P},\ k \in \mathcal{N}_t \tag{2.2.6}$$

where the allowable ramp $\overline{\Delta P}$ is often tuned to prevent premature damage to the battery. The residual demand d_k requested from the utility is:

$$d_k = \hat{L}_k - P_k + \hat{\alpha}_k F_k,\ k \in \mathcal{N}_t \tag{2.2.7}$$

We assume that the ISO does not allow the battery to sell back electricity. This is modeled

by using the constraint:

$$P_k + F_k \leq \hat{L}_k,\ k \in \mathcal{N}_t \tag{2.2.8}$$

This constraint can also be written as $\hat{L}_k - P_k - F_k \geq 0$. Consequently, when the ISO fully uses the FR capacity we have that $\alpha_k = -1$ and the constraint implies that $d_k \geq 0$ because $d_k = \hat{L}_k - P_k - \alpha_k F_k$. This guarantees that the battery does not sell excess power.

We use the following peak demand *carryover* constraint:

$$\max_{k \in \mathcal{N}_t} d_k \geq \widehat{D}_t \tag{2.2.9}$$

Finally, we impose the following basic capacity bounds:

$$0 \leq E_k \leq \overline{E},\ k \in \mathcal{N}_t \tag{2.2.10a}$$

$$-\underline{P} \leq P_k \leq \overline{P},\ k \in \mathcal{N}_t \tag{2.2.10b}$$

$$0 \leq F_k \leq \overline{P},\ k \in \mathcal{N}_t \tag{2.2.10c}$$

We denote the optimization problem solved at time t as $\mathcal{P}_t(\hat{L}_{\mathcal{N}_t}, \hat{\alpha}_{\mathcal{N}_t}, \hat{\pi}_{\mathcal{N}_t}, E_t, P_t, \widehat{D}_t)$ where the arguments are the *input data* needed to solve the problem. When convenient, we drop the dependence on the forecast prices, loads, and FR signals and use the compact form $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$.

We note that the deterministic formulation uses forecasts of uncertain prices, loads, and FR dispatch signals. To represent uncertainty in these variables we model them as random variables with realizations (realizations) denoted by $\xi \in \bar{\Xi}$ (where Ξ is the scenario set). In particular, at time t, a realization of the uncertain load, prices, and FR signals are denoted as $L_t(\xi)$, $\pi_t(\xi)$ and $\alpha_t(\xi)$, respectively.

We can now summarize the receding horizon scheme run over the monthly planning period $\mathcal{M}$ as follows:

- START at $t = 0$ with E_0 and $\widehat{D}_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:

- Solve $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$ by using the forecast prices $\hat{\pi}_{\mathcal{N}_t}$, loads $\hat{L}_{\mathcal{N}_t}$, and FR signals $\hat{\alpha}_{\mathcal{N}_t}$ to obtain commitments P_{t+1} and F_{t+1}. Implement these decisions over $(t, t+1)$ and update the new battery SOC to $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$ using the true FR signal $\alpha_{t+1}(\xi)$.

- Using the true load $L_{t+1}(\xi)$ and FR signal $\alpha_{t+1}(\xi)$, compute the true demand $d_{t+1}(\xi) = L_{t+1} - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$. Update the carryover peak as $\widehat{D}_{t+1}(\xi) = \max\{\widehat{D}_t(\xi), d_{t+1}(\xi)\}$.

- Set $t \leftarrow t+1$. If $t = M$, STOP, otherwise RETURN to second step.

2.3 Stochastic MPC Framework

A limitation of deterministic MPC is that it only uses forecast information of the loads and prices but neglects their uncertainty when making commitments. This can introduce significant inefficiencies (e.g., if the realized true load is significantly different than the forecasted one). We now derive a stochastic MPC framework that factors in uncertainties on loads and prices explicitly by formulating a two-stage stochastic program that is solved over a receding horizon.

We define the net power P_{t+1} and FR commitments F_{t+1} over the next immediate trading hour as *first-stage decisions* (here-and-now) that need to be made prior to observing uncertainty (i.e., when the market is settled). The power and regulation capacity trajectories $P_k(\xi)$ and $F_k(\xi)$ for all $k \in \mathcal{N}_t \setminus \{t+1\}$ are modeled as second-stage or *recourse decisions* that can be corrected once uncertainty reveals. We note that the residual and peak demands are also recourse variables that we express as $d_{\mathcal{N}_t}(\xi)$ and $D_t(\xi) = \max_{k \in \mathcal{N}_t} d_k(\xi)$. We highlight that the SOC at time $t+1$ only depends on the previous storage level E_t and on the first-stage variable P_{t+1}. Consequently, E_{t+1} is also a first-stage variable. The rest of the trajectory $E_k(\xi)$ for $k \in \mathcal{N}_t \setminus \{t+1\}$ are second-stage variables because $P_k(\xi)$ are second-stage variables for $k \in \mathcal{N}_t \setminus \{t+1\}$.

2.3.1 *Objective Function*

In the stochastic formulation we maximize the expected revenue over horizon $\mathcal{N}_t$:

$$\mathbb{E}\left[\sum_{k\in\mathcal{N}_t}\pi^e_k(\xi)(P_k(\xi)-\alpha_k(\xi)F_k(\xi))+\sum_{k\in\mathcal{N}_t}\pi^f_k(\xi)F_k(\xi)-\frac{\pi^D}{\sigma_t}\max_{k\in\mathcal{N}_t}d_k(\xi)\right] \tag{2.3.11}$$

2.3.2 *Constraints*

The constraints of the stochastic formulation are essentially the same as those of the deterministic counterpart but are replicated for every realization $\xi \in \bar{\Xi}$ where $\bar{\Xi} \subseteq \Xi$. Here, we state them explicitly for completeness. The limits on the net battery discharge are:

$$P_k(\xi)+F_k(\xi)\leq \overline{P},\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.12a}$$

$$P_k(\xi)-F_k(\xi)\geq -\underline{P},\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.12b}$$

The dynamics on the storage level are given by:

$$E_k(\xi) = E_{k-1}(\xi)-P_k(\xi)+\alpha_k(\xi)F_k(\xi),\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.13}$$

The cushion storage constraints are given by:

$$\rho F_k(\xi)\leq E_{k-1}(\xi)\leq \overline{E}-\rho F_k(\xi),\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.14a}$$

$$\rho F_k(\xi)\leq E_k(\xi)\leq \overline{E}-\rho F_k(\xi),\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.14b}$$

The ramping limits on the net power are given by:

$$-\overline{\Delta P}\leq P_k(\xi)-P_{k-1}(\xi)\leq \overline{\Delta P},\ k\in\mathcal{N}_t,\xi\in\bar{\Xi} \tag{2.3.15}$$

The residual demand is given by:

$$d_k(\xi) = L_k(\xi) - P_k(\xi) + \alpha_k(\xi)F_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{2.3.16}$$

The no-sellback restriction to the ISO is given by:

$$P_k(\xi) + F_k(\xi) \leq L_k(\xi),\ k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{2.3.17}$$

The bounds on the variables are given by:

$$0 \leq E_k(\xi) \leq \overline{E}, k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{2.3.18a}$$

$$-\underline{P} \leq P_k(\xi) \leq \overline{P}, k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{2.3.18b}$$

$$0 \leq F_k(\xi) \leq \overline{P}, k \in \mathcal{N}_t, \xi \in \bar{\Xi} \tag{2.3.18c}$$

The peak demand must satisfy:

$$\max_{k \in \mathcal{N}_t} d_k(\xi) \geq \widehat{D}_t,\ \xi \in \bar{\Xi} \tag{2.3.19}$$

Finally, we enforce the fact that P_{t+1} and F_{t+1} are first-stage variables by using the non-anticipativity constraints:

$$P_{t+1}(\xi) = \mathbb{E}\left[P_{t+1}(\xi)\right],\ \xi \in \bar{\Xi} \tag{2.3.20a}$$

$$F_{t+1}(\xi) = \mathbb{E}\left[F_{t+1}(\xi)\right],\ \xi \in \bar{\Xi} \tag{2.3.20b}$$

Figure 2.4 sketches the implementation of the stochastic receding horizon scheme. At each time step, we solve an N-hour horizon two-stage stochastic program and obtain a single value for the commitments and of the state of charge (SOC) at the immediate next hour (shown in bold black lines and dots) and multiple realizations for the state trajectory for the remaining hours (shown in colored dashed lines and dots). The implementation is

essentially the same as the deterministic counterpart but with the important observation that the problem at time t uses a forecast and realizations for the uncertain variables that we denote as $L_{\mathcal{N}_t}(\xi)$, $\alpha_{\mathcal{N}_t}(\xi)$, $\pi^f_{\mathcal{N}_t}(\xi)$, and $\pi^e_{\mathcal{N}_t}(\xi)$. So, we can denote the problem that is solved at time t as $\mathcal{P}_t(L_{\mathcal{N}_t}(\xi), \alpha_{\mathcal{N}_t}(\xi), \pi^f_{\mathcal{N}_t}(\xi), \pi^e_{\mathcal{N}_t}(\xi), E_t, P_t, \widehat{D}_t)$. For convenience, we simplify the notation to $\mathcal{P}_t(E_t, P_t, \widehat{D}_t)$. From the solution to this problem we obtain the commitments P_{t+1} and F_{t+1}. The actual residual demand, however, depends on the actual load realized and given by $d_{t+1} = L_{t+1} - P_{t+1} - \alpha_{t+1}F_{t+1}$ where L_{t+1} and α_{t+1} are the actual (real) values of the load and dispatch signal.

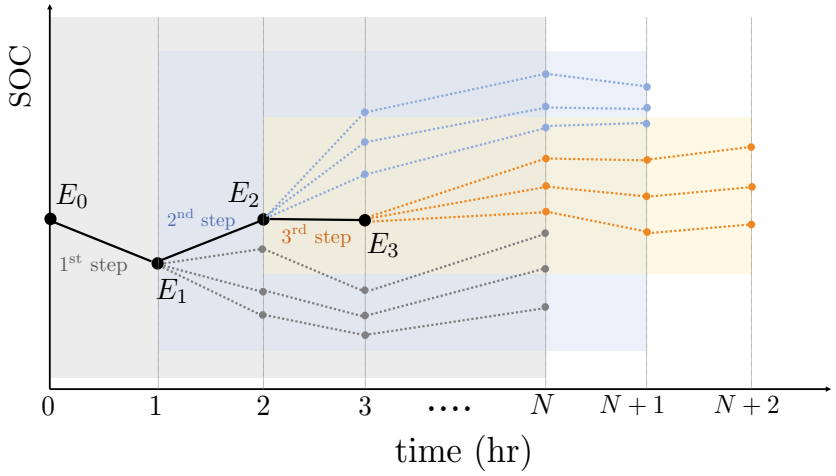

Figure 2.4: Sketch of stochastic MPC scheme.

We can now summarize the stochastic MPC scheme run over a month period $\mathcal{M}$ as follows:

- START at $t = 0$ with E_0 and $\widehat{D}_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:
- Solve $\mathcal{P}_t\left(E_t, P_t, \widehat{D}_t\right)$ by using the forecast realizations $\pi_{\mathcal{N}_t}(\xi)$, loads $L_{\mathcal{N}_t}(\xi)$, and FR signals $\alpha_{\mathcal{N}_t}(\xi)$ to obtain commitments P_{t+1} and F_{t+1}. Implement these decisions over $(t, t+1)$ and update the battery SOC $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}F_{t+1}$ using the *realized* FR signal α_{t+1}.
- Using the *actual* realized load L_{t+1} and FR signal α_{t+1} compute the actual net demand $d_{t+1} = L_{t+1} - P_{t+1} + \alpha_{t+1}F_{t+1}$. Update peak carryover demand as $\widehat{D}_{t+1} =$

$\max\{\widehat{D}_t, d_{t+1}\}$.

- Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN to second step.

We highlight that we impose a lower bound for the peak demand $D_t = \max_{k\in\mathcal{N}_t} d_k$ given by $\widehat{D}_t$, which is a result of carrying over the observed peak until time t in the receding horizon scheme. The need for such a lower bound can be observed by noticing that the peak demand over the entire month horizon $\mathcal{M}$ can be computed as a recursion of the form $\widehat{D}_{k+1} = \max(\widehat{D}_k, d_{k+1}), k \in \mathcal{M}$ with $\widehat{D}_0 = 0$. Consequently, at a given time t, the peak demand D_t over the horizon $\mathcal{N}_t$ must be at least as much as $\widehat{D}_t$, and thus $D_t \geq \widehat{D}_t$. This implies that $\max_{k\in\mathcal{N}_t} d_k \geq \widehat{D}_t$ must hold. This *carryover* constraint indicates that the peak demand is, in fact, a state of the system, and thus must be propagated forward in time (as in the case of the battery SOC).

2.4 Benchmarking Stochastic MPC

In this section, we compare the performance of the proposed stochastic MPC framework against that of deterministic MPC and perfect information MPC. We also describe a methodology to generate realizations from historical data.

2.4.1 *Performance Metrics*

To quantify the impact of forecast errors on the performance of *deterministic MPC*, we use the commitment trajectories $P_{\mathcal{M}}$ and $F_{\mathcal{M}}$ collected over the month and different *realizations* of the loads $L_{\mathcal{M}}(\xi)$, prices $\pi_{\mathcal{M}}(\xi)$, and FR signals $\alpha_{\mathcal{M}}(\xi)$, to compute the actual monthly net revenue. Since the revenue is negative, we find it more convenient to use the total cost (the negative net revenue). This is a random variable with realizations:

$$\Phi^{det}(\xi) = \sum_{t\in\mathcal{M}} \pi_t^e(\xi)(\alpha_t(\xi)F_t - P_t) - \sum_{t\in\mathcal{M}} \pi_t^f(\xi)F_t + \pi^D \max_{t\in\mathcal{M}} d_t(\xi), \quad \xi \in \Xi \tag{2.4.21}$$

where $d_t(\xi) = L_k(\xi) - P_t - \alpha_t(\xi)F_t$.

From the commitment policies $P_{\mathcal{M}}$, $F_{\mathcal{M}}$ obtained under the *stochastic MPC scheme* and the *actual* realizations of the loads, prices, and FR signals, we can compute the accumulated revenue for every realization using (2.4.21) as well (i.e., we use the same realizations but different commitment policies). The net revenue under the stochastic policy is denoted as $\Phi^{sto}(\xi)$. We note that the revenue of stochastic and deterministic MPC is computed for *every realization* of the uncertain variables $\xi \in \Xi$. This accounts for the fact that, in practice, the stochastic MPC problem cannot be solved for every possible realization of uncertainty (i.e., we often have that $|\bar{\Xi}| \ll |\Xi|$).

To evaluate the *ideal* performance of MPC (we call this *perfect information MPC*), we compute commitment policies at every time t for discharge/charge power and FR capacity for *every realization* $\xi \in \Xi$ of the loads, prices, and FR signals. The monthly total cost under this scheme is given by:

$$\Phi^{perf}(\xi) = \sum_{t\in\mathcal{M}} \pi_t^e(\alpha_t(\xi)F_t(\xi) - P_t(\xi)) - \sum_{t\in\mathcal{M}} \pi_t^f(\xi)F_t(\xi) + \pi^D \max_{t\in\mathcal{M}} d_t(\xi), \quad \xi \in \Xi. \tag{2.4.22}$$

Note that the decisions are a function of the realizations, which highlights the fact that we compute policies by perfectly anticipating uncertainty. We also highlight that these perfect information policies are not implementable in practice and therefore are only used as a benchmark.

In stochastic programming studies, one is often interested in evaluating the expected value of the stochastic solution (which is defined as $\mathbb{E}[\Phi^{stoch}(\xi) - \Phi^{det}(\xi)]$ and denoted as VSS). It is also often convenient to assess the value of having perfect information, which we define as $\mathbb{E}[\Phi^{det}(\xi) - \Phi^{perf}(\xi)]$ and which is denoted as VPI (Birge and Louveaux, 2011). The first metric evaluates the benefit of stochastic over deterministic MPC (which neglects uncertainty at the moment of decision) while the second metric evaluates the impact of not being able to predict market and load conditions perfectly. We highlight that our definition

of VPI uses deterministic MPC as a basis, instead of stochastic MPC, as is usually done in the literature. With this, we seek to express all our metrics relative to deterministic MPC, which is the standard approach used in industrial applications.

Of particular interest in our context is a metric that we call the *expected value of the battery*. To compute this value, we evaluate the performance of the system under the assumption that no battery has been installed (and thus FR revenue and peak shaving capabilities are not available). The total cost thus reduces to the demand charge:

$$\Phi^{nobat}(\xi) = \pi^D \max_{t \in \mathcal{M}} L_t(\xi). \tag{2.4.23}$$

The ideal expected value of the battery (the ideal expected savings) is thus defined as $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{perf}(\xi)]$. The value of the battery under stochastic MPC is $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{stoch}(\xi)]$ and under deterministic MPC is $\mathbb{E}[\Phi^{nobat}(\xi) - \Phi^{det}(\xi)]$. As we will see, these metrics can be used to isolate the effect of the battery, which is usually significantly smaller than the total cost. If the battery provides value, we have that these quantities are positive. Another important quantity that can be derived directly from the value of the battery for each policy is the *payback period*, which is simply the ratio of the investment cost and the value of the battery (i.e., the expected savings).

We highlight that, as discussed in Section 2.3.1, the cost terms $\pi_t^e(\xi)L_t(\xi)$ are not considered in the evaluation of the net revenues of the different policies because this is a large and constant term that is common across policies. We also note that the monthly revenues do not use demand charge discounting (as done in MPC objective functions). This is because discounting is a heuristic strategy to deal with the fact that, in general, one cannot solve MPC problems with extremely long horizons.

2.4.2 *Scenario Generation from Limited Historical Data*

One often is interested in using historical data to model the uncertainty of random variables and generate scenarios (realizations). In our context, we seek to generate scenarios

by sampling an empirical multivariable Gaussian distribution that captures short-term and long-term time correlations. Unfortunately, historical data is often insufficient to obtain a sample covariance. For instance, if we seek to capture daily and weekly correlations, we can only obtain 52 samples from a full year of data but the associated sample covariance has a dimension of $24 \times 7 = 168$. Consequently, estimating a sample covariance matrix from the historical profiles is not possible (the covariance matrix does not have full rank and thus cannot be used for sampling). To address this issue, we use the Ledoit-Wolf covariance estimator for large dimensional vectors proposed in Ledoit and Wolf (2004). The Ledoit-Wolf covariance estimator is both well-conditioned and converges to the sample covariance matrix asymptotically. In practice, it has been observed that a handful of observations are sufficient to obtain a good covariance approximation. The estimator and its convergence properties are described in detail in Ledoit and Wolf (2004), here we focus only on its essential features.

In the following description, we consider the norm of a $p \times p$ matrix to be the a scaled version of the Frobenius norm $||A||^2 = \text{Tr}(AA^T)/p$ and the associated inner product to be $\langle A_1, A_2 \rangle = \text{Tr}(A_1 A_2^T)/p$. The advantage of dividing the trace by the dimension p is that the norm of the identity matrix is one. Assume now that we have n independent and identically distributed (i.i.d.) observations of a p-dimensional multivariate random variable X with mean 0 and covariance Σ_n, where $p \gg n$ (in our case $n = 52$ and $p = 168$). Let X_n denote a $p \times n$ matrix of the n realizations of the x vector. Our goal is to find the linear combination $\Sigma_n^* = \rho_1 I_n + \rho_2 S_n$ of the identity matrix I_n and the sample covariance matrix $S_n = X_n X_n^T/n$. Here, $\rho_1, \rho_2 \geq 0$ and we define the expected quadratic loss $\mathbb{E}[||\Sigma_n^* - \Sigma_n||]$ where Σ_n is the true (unobservable) covariance matrix. We also define the following scalar quantities, which play a central role in the analysis: $\mu_n = \langle \Sigma_n, I_n \rangle, \alpha_n^2 = ||\Sigma_n - \mu_n I_n||^2, \beta_n^2 = \mathbb{E}[||S_n - \Sigma_n||^2]$, and $\delta_n^2 = \mathbb{E}[||S_n - \mu_n I_n||^2]$. It is not necessary to assume that the random variables follow a specific distribution, but we need to assume that they have finite moments (so that β_n^2 and δ_n^2 are finite). Using these definitions, it can be shown that $\mathbb{E}[S_n] = \Sigma_n$ and that $\alpha_n^2 + \beta_n^2 = \delta_n^2$ hold.

We obtain the optimal combination $\Sigma_n^* = \rho_1 I_n + \rho_2 S_n$ by solving the quadratic program:

$$\min_{\rho_1,\rho_2} \quad \mathbb{E}[||\Sigma_n^* - \Sigma_n||^2] \tag{2.4.24a}$$

$$\text{s.t.} \quad \Sigma_n^* = \rho_1 I_n + \rho_2 S_n \tag{2.4.24b}$$

The solution of this quadratic program is:

$$\rho_1 = \frac{\beta_n^2}{\delta_n^2}\mu_n, \ \rho_2 = \frac{\alpha_n^2}{\delta_n^2}. \tag{2.4.25}$$

The optimal covariance matrix Σ_n^* obtained by substituting (2.4.25) into (2.4.24b), however, is not a *bona fide* estimator because it requires hindsight knowledge of the scalar functions of the true (but unobservable) covariance matrix Σ_n ($\mu_n, \alpha_n^2, \beta_n^2$, and δ_n^2). This problem is addressed by obtaining an alternative consistent estimator for Σ_n^*. To do this, we note that computing Σ_n^* does not require knowledge of the whole matrix Σ_n directly, but only of the scalar functions $\mu_n, \alpha_n^2, \beta_n^2$, and δ_n^2. We thus define consistent estimators of $\mu_n, \alpha_n^2, \beta_n^2$ and δ_n^2 that converge to the true ones for large n. We define $m_n = \langle S_n, I_n \rangle$, $d_n^2 = ||S_n - m_n I_n||^2$, $\bar{b}_n^2 = \frac{1}{n^2}\sum_{k=1}^n \left|\left|x_k x_k^T - S_n\right|\right|^2$, $b_n^2 = \min(\bar{b}_n^2, d_n^2)$, and $a_n^2 = d_n^2 - b_n^2$. It is possible to prove that $\mathbb{E}[m_n] = \mu_n$, and that $m_n - \mu_n$, $m_n^2 - \mu_n^2$, $d_n^2 - \delta_n^2$, $\bar{b}_n^2 - \beta_n^2$ and $b_n^2 - \beta_n^2$, $a_n^2 - \alpha_n^2$ all converge to 0 as n increases. The next step is to replace the unobservable scalars in (2.4.24b) with consistent estimators. This yields the *bona fide* estimator:

$$\tilde{\Sigma}_n = \frac{b_n^2}{d_n^2} m_n I_n + \frac{a_n^2}{d_n^2} S_n \tag{2.4.26}$$

We use this covariance estimator $\tilde{\Sigma}_n$ to generate samples for uncertain variables in our benchmark studies.

2.4.3 *Computational Experiments*

We evaluate the performance of the stochastic MPC framework proposed by simulating a utility-scale stationary battery that has a capacity 0.5 MWh with a rated power of 1 MW for both charge and discharge. A ramping limit of 0.5 MW/hr was used in the simulations. We generate realizations for energy prices and FR prices and signals using data from PJM. Historical load data for a typical university campus is used to generate realizations for the load. We consider cases with MPC prediction horizon lengths of $N = 24$ (1-day horizon), $N = 72$ (3-day horizon), and $N = 168$ (7-day horizon) to study the effect of the horizon length on economic performance. The horizon lengths that we consider are short compared to the demand charge period (one month). Consequently, we also consider different strategies to estimate the *discounting factor* σ_t. In the first case we consider a constant discounting factor of $\sigma_t = 1$, which means that no discounting of the long-term demand

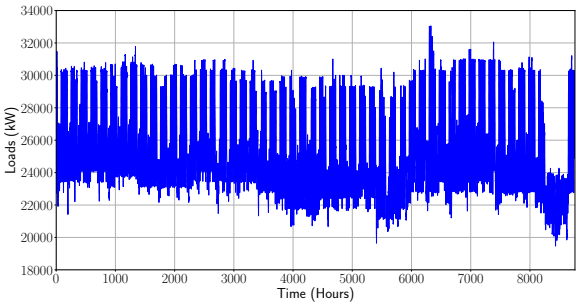

(a) Full year

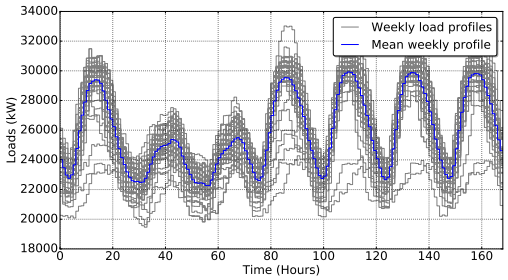

(b) Weekly load profiles used for Ledoit-Wolf estimator

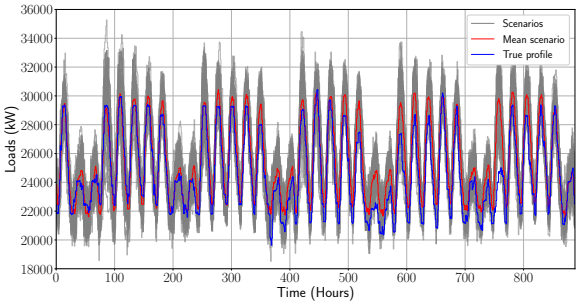

(c) Generated weekly load realizations using Ledoit-Wolf estimator

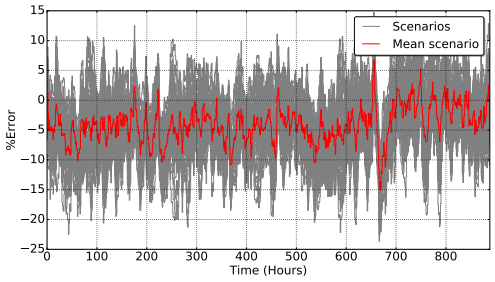

(d) Prediction errors for mean and scenario profiles.

Figure 2.5: Historical loads and realizations

charge rate is used over the short horizon of length N. In the second case we use a discounting $\sigma_t = M/N$. Here, we recall that π_d is the demand charge price over a month and thus the discounting factor $\sigma_t = M/N$ scales down the price when the horizon is shorter than a month. For instance, if the horizon is one day, we scale down the demand charge by a factor of $\sigma_t = 30 \times 24/24 = 30$ and if the horizon is one week, we scale down the demand charge by a factor of $\sigma_t = 30 \times 24/7 \times 24 = 30/7$.

We perform receding-horizon simulations for a period of one month. Load realizations for this period are generated from one year of historical data using the Ledoit-Wolf covariance estimator. The one-year load profile of a typical large university campus is shown in Figure 2.5. Figure 2.5b shows the weekly campus load profiles that were used to construct the covariance estimator while the weekly load realizations generated with the Ledoit-Wolf estimator are shown in Figure 2.5c. The prediction errors for the realizations and for the mean profile (the most likely profile) are shown in Figure 2.5d. We highlight that the

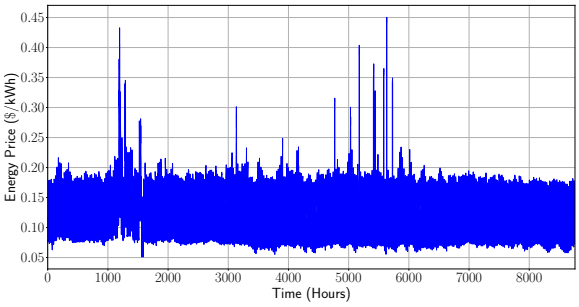

(a) Full year

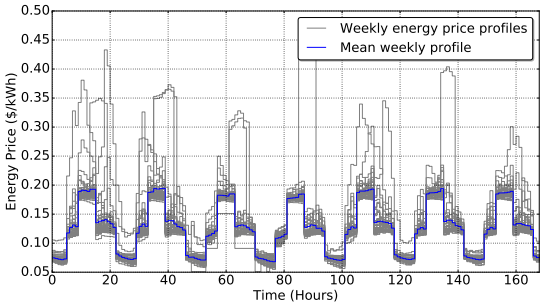

(b) Weekly energy price profiles used in Ledoit-Wolf estimator

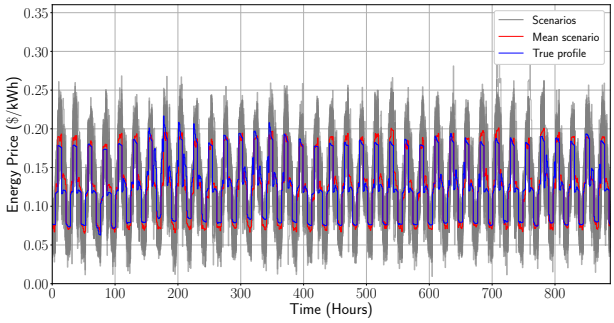

(c) Generated energy price realizations using Ledoit-Wolf estimator

Figure 2.6: Historical energy prices and realizations

errors remain in the range of $[-20\%, 10\%]$ (relative to the true load) and emphasize that the error of the mean profile is in the range of -15% to +5%. This indicates that the proposed load scenario generation approach provides reasonable estimates and is mostly due to the fact that the load exhibits strong periodicity. Moreover, we note that the estimator preserves the time correlation structure of the load. For some load realizations, we see a significant deviation from the trend, due to seasonal effects (on a year scale). This could potentially be resolved by extending the load realizations to capture monthly correlations but this will increase historical data needs and deteriorate the performance of the Ledoit-Wolf estimator (i.e., for a one-year horizon, we only have 12 realizations while the random variable has a dimension of 720). We will explore this more extreme case in future work.

The one-year energy price profile obtained from PJM is shown in Figure 2.6a. We observe that the profile also exhibits strong periodicity. We also use a Ledoit-Wolf covariance

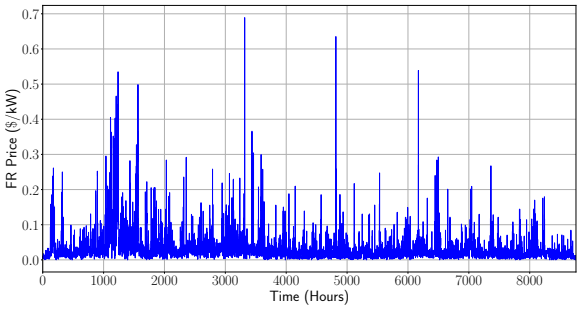

(a) Full year

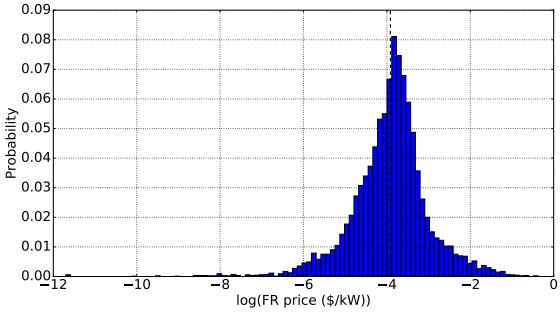

(b) Distribution of logarithm of FR prices for one year

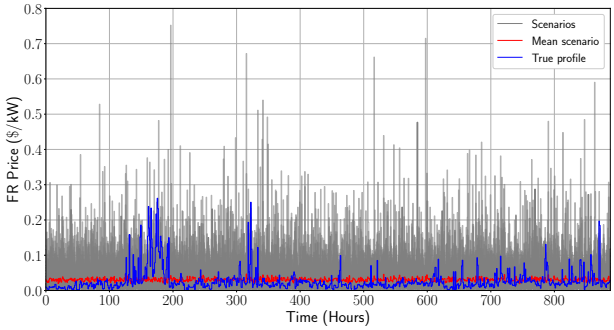

(c) Generated FR price realizations

Figure 2.7: Historical FR prices and realizations

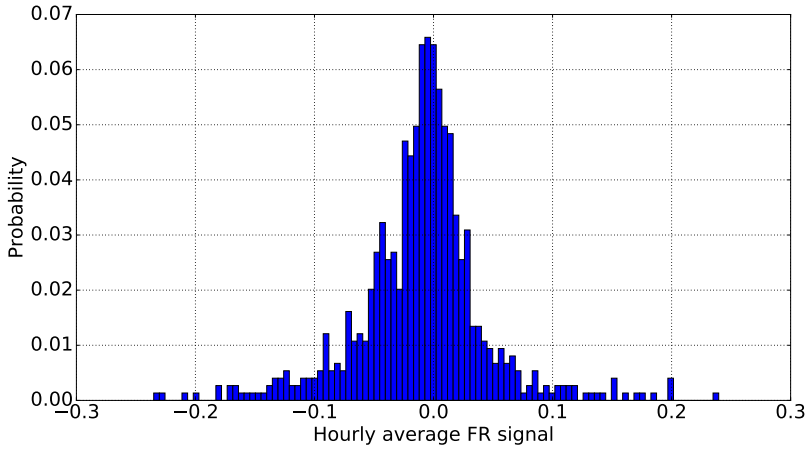

(a) Distribution of hourly average FR dispatch signal

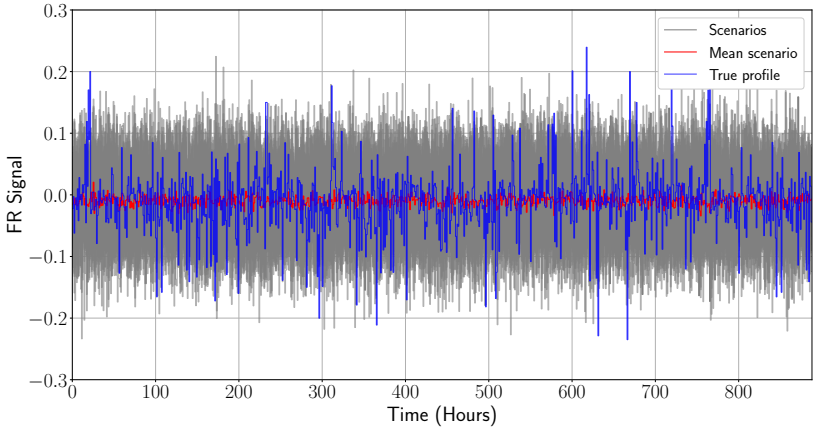

(b) Generated hourly average FR signal realizations

Figure 2.8: Historical hourly FR signals and realizations

estimator to generate realizations. Figure 2.6b shows the weekly energy price profiles that were used to construct the covariance estimator while the weekly energy price realizations generated with the Ledoit-Wolf estimator are shown in Figure 2.6c.

The FR price data obtained from PJM is shown Figure 2.7a. We observe that, compared to the energy price, the FR price profile exhibits significantly less periodicity and time correlation. From Figures and 2.7c and 2.7b we observe that the FR price follows a log-normal distribution but with high variance. We use this distribution to generate FR price realizations, which are shown in Figure 2.7c. We note that that the significant conservatism is expected with respect to this variable.

To generate realizations of the hourly average FR dispatch signal, we note from Figure 2.8a that the signal approximately follows a Gaussian distribution, which also indicates that there is small temporal correlation. Figure 2.8b shows the realizations generated by sampling from the empirical Gaussian.

For the simulation of deterministic MPC, we use the mean (most likely) profiles of the random variables as forecasts. We evaluate the policy in $|\Xi| = 1000$ realizations to evaluate the distribution of the monthly cost. In the stochastic MPC scheme, we solve a two stochastic stochastic program containing $|\bar{X}| = 50$ realizations to compute the planning policy and evaluate the policy in the same 1000 realizations used for the deterministic and perfect information schemes. To simulate the perfect information MPC scheme, we treat the realizations of uncertain parameters as the actual ones and optimize over each one individually (this is equivalent to dropping the non-anticipativity constraints) (Birge and Louveaux, 2011). All optimization problems in the MPC simulations were implemented in `Julia` using the modeling package `JuMP` and solved using `Gurobi 6.5`.

Total Costs and Peak Demands

In Figure 2.9, we compare the distribution of the total costs (negative net revenues) for the 1000 realizations evaluated. As can be seen, the *tail costs* in the deterministic policy are significantly more pronounced, indicating that a larger number of realizations experience high costs. The stochastic MPC scheme mitigates these extreme costs and also has a *mean cost shifted towards the leftt*. The cost distribution for perfect information MPC is clearly superior but we emphasize that this approach is not implementable and is only presented to illustrate ideal behavior.

From Figure 2.9 we can also observe that the distribution of the total costs for all MPC schemes shift towards the left as we increase the prediction horizon. This indicates that extending the horizon is beneficial. In particular, strong differences in the distributions are observed for 1-day and 3-day horizons while the differences between the 3-day and 7-day policies are less pronounced. This indicates that the prediction horizon plays a key

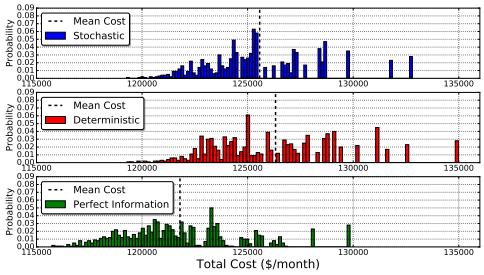

(a) 1-day horizon, undiscounted demand

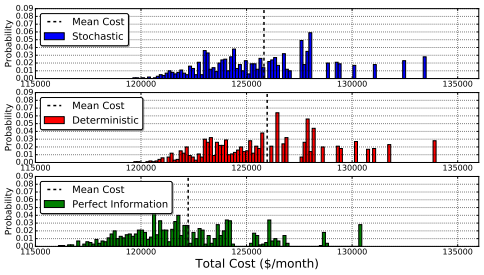

(b) 1-day horizon, discounted demand

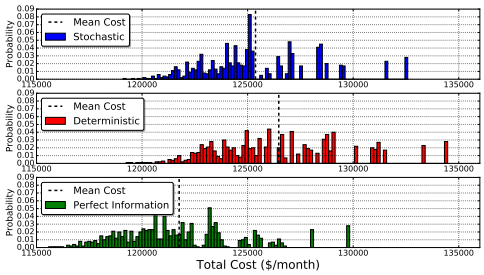

(c) 3-day horizon, undiscounted demand

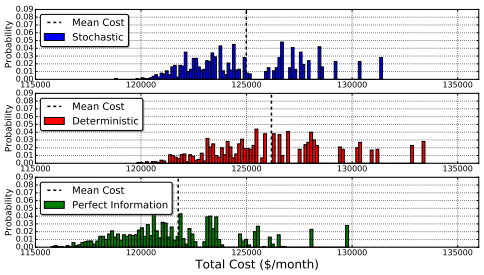

(d) 3-day horizon, discounted demand

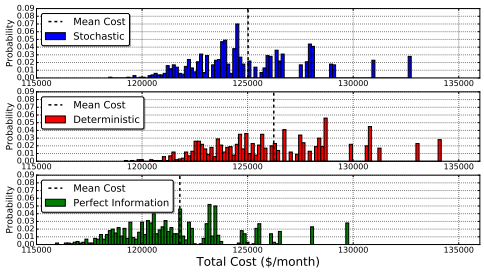

(e) 7-day horizon, undiscounted demand

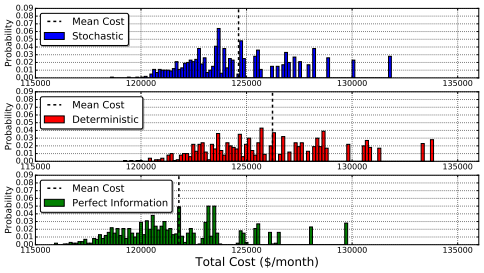

(f) 7-day horizon, discounted demand

Figure 2.9: Distribution of total monthly costs (1000 realizations)

role, as is typically observed in MPC applications, *but using a very long horizon does not seem necessary*. From Figure 2.9 we can also observe that, when the demand charge is discounted as a factor of the horizon length, the cost distribution shifts towards the left. This is particularly evident in the 1-day horizon cases. This shows that it is important to properly discount the peak demand when we solve the optimization problem for shorter horizons. This is because an *undiscounted demand charge will force the MPC system to operate more conservatively* (i.e., by trying to mitigate the peak demand over the prediction MPC horizon at every solution time t).

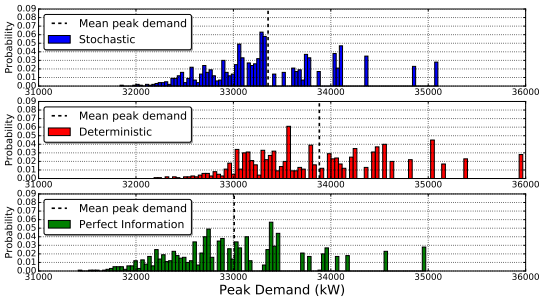

(a) 1-day horizon, undiscounted demand

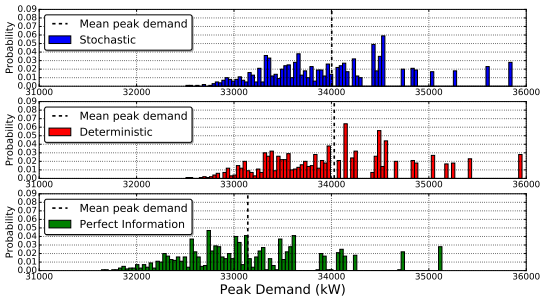

(b) 1-day horizon, discounted demand

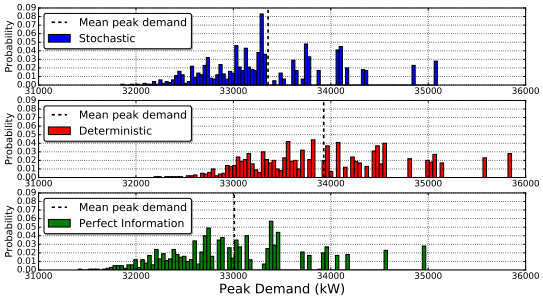

(c) 3-day horizon, undiscounted demand

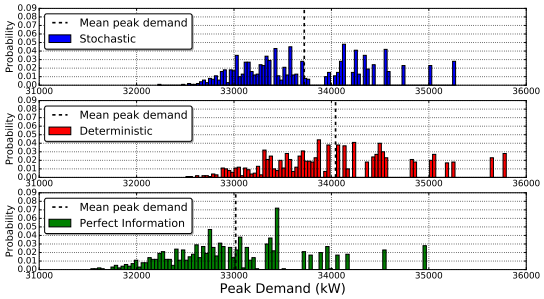

(d) 3-day horizon, discounted demand

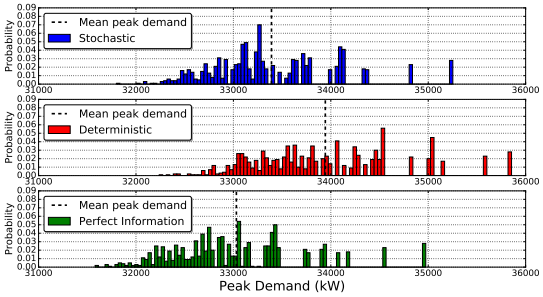

(e) 7-day horizon, undiscounted demand

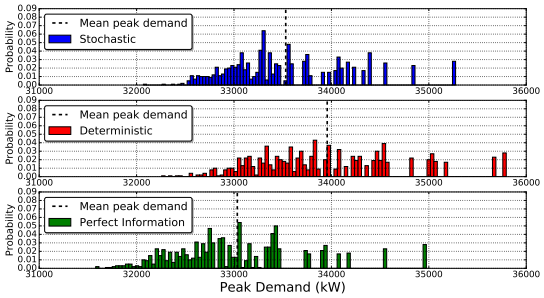

(f) 7-day horizon, discounted demand

Figure 2.10: Distribution of peak demands over a month (1000 realizations)

From Figure 2.10, we also observe that the peak demand is distributed towards lower values for longer horizons. Again, this highlights the importance of choosing a sufficiently long horizon. The stochastic MPC scheme is better able to mitigate the peaks compared to deterministic MPC. This can also be seen from the random realization of the load profile show in in Figure 2.11. Figure 2.10 also shows that, when the peak demand is discounted, the peak demands are distributed towards lower values. These observations illustrate that using incorrect discounting factors and insufficient horizon lengths can undermine the

value of the stochastic solution.

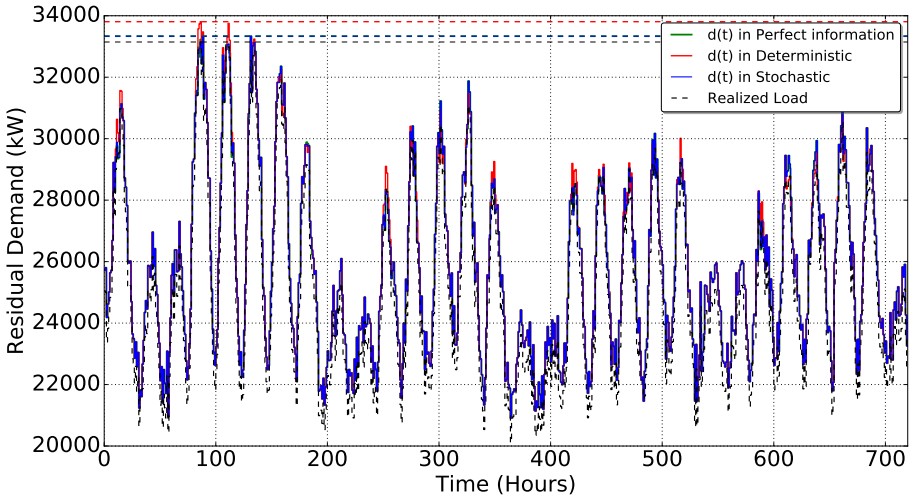

Figure 2.11: Peak shaving performance for stochastic, deterministic, and perfect information MPC in a random load profile (7-day horizon, discounted demand)

FR Commitments and SOC Policies

In Figure 2.12, we compare the *regulation commitment levels* of the different schemes. We observe that the FR participation levels are consistent in all policies, with most of them

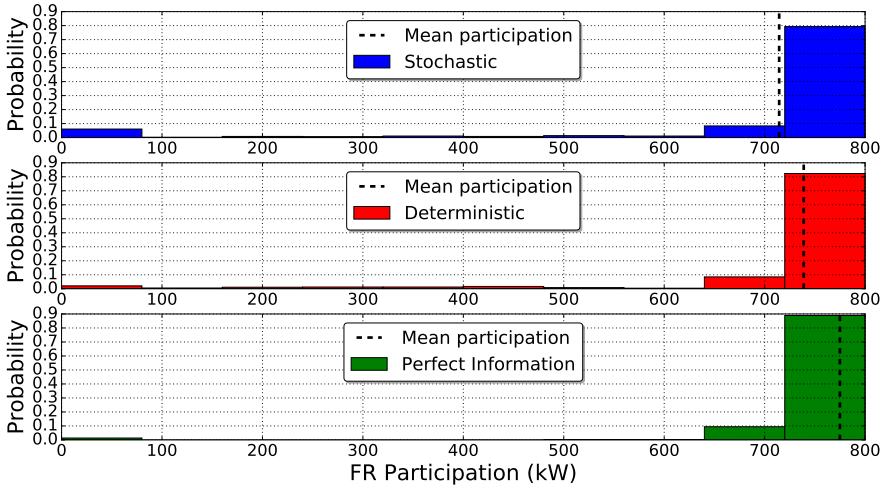

Figure 2.12: Distributions of FR commitments (7-day horizon, discounted demand)

clustering in the range of 700-800 kW. We also observe that the deterministic policy is slightly more aggressive in allocating FR capacities than the stochastic policy (i.e., the mean of the FR commitment is shifted towards the right). This comes at the expense of higher peak demand (as seen in Figure 2.10). This highlights that the stochastic policy can better manage stored energy to protect against high peak demands. Figure 2.13 compares the *dynamics of the SOC* (for a random scenario) for the case study with a 7-day horizon and with a discounted demand charge cost. Here, we can observe that the deterministic policy exhibits abrupt changes in stored energy more frequently than the stochastic and perfect information schemes. This reflects the fact that the deterministic policy overcommits in FR capacity. The SOC profiles of stochastic and perfect information MPC are quite similar and smoother. This indicates that stochastic MPC performance is close to that of perfect information MPC.

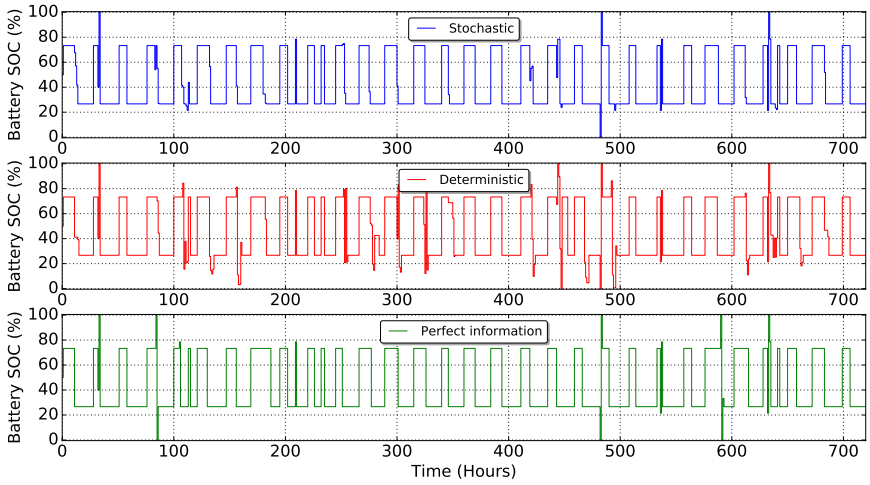

Figure 2.13: SOC policy in stochastic, deterministic and perfect information MPC (7-day horizon, discounted demand charge)

Value of the Battery

Table 2.1 reports the total cost and its components for all the case studies. We observe a general trend of higher expected total costs with shorter horizon lengths and without

discounted factors for all the schemes. Moreover, we observe that stochastic MPC outperforms deterministic MPC. We also highlight that the stochastic MPC policy is computed by using only 50 realizations, which seems sufficient to outperform deterministic MPC. In Table 2.2, we compare the value of the stochastic solution (VSS) and the value of perfect information (VPI). In our studies, we see that VSS can reach levels of $1,600 per month while VPI can reach up to $4,700 per month. Both of these quantities are small compared to the total cost (which reaches levels of $120,000 per month). This is because the load of the university campus (and thus the demand charge) is significantly larger than the battery capacity, which is usually the case in demand-side management applications. As a result, the VSS and VPI metrics are not particularly informative on the added value of the battery management strategy.

Table 2.3 summarizes the results for the *value of the battery* under different schemes for a 7-day horizon and discounted demand charge. The expected savings obtained by installing the battery and operating it under perfect information (*ideal value of the battery*) reach $16,430 per month (approximately $197,160 per year). This reveals that the battery is indeed valuable. From Table 2.2 we observe that FR revenue reaches levels of $14,000 per month (approximately $168,000 per year). From Table 2.3 we observe that stochastic MPC can recover *83%* of the ideal value of the battery while deterministic MPC can only recover 73%. Assuming an investment cost of 1.0 M$, the payback period of the battery operated under perfect information MPC is 5.0 years, under stochastic MPC is 6.1 years, and under deterministic MPC is 6.94 years. Stochastic MPC achieves a relative improvement in the payback period achieved under deterministic MPC by 12.1%. By comparing the payback period under perfect information and deterministic MPC, we can also conclude that perfect information achieves a relative improvement of 27.9%. It is rather remarkable to see that, by improving the battery management strategy, one can significantly influence the value of the battery and the associated payback period. From these results, we also see that the value of the battery and payback period are more informative benchmark metrics than the more generic VSS and VPI metrics.

Table 2.1: Expected costs for MPC schemes. TC= Total cost, DC=Demand charge, FR=frequency regulation, and E=Energy.

Horizon	**7-day**					
σ_t	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
Cost Item ($/month)	**Perfect Information**		**Stochastic**		**Deterministic**	
TC	121,788.9	121,794.4	125,022.5	124,621.7	126,237.8	126,229.9
DC	137,077.0	137,090.5	138,572.9	139,153.3	140,865.7	140,920.5
FR	-14,812.9	-14,820.8	-13,073.7	-14,049.7	-14,116.8	-14,180.9
E	-475.1	-475.2	-476.6	-481.9	-511.1	-509.7
Horizon	**3-day**					
σ_t	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
Cost Item ($/month)	**Perfect Information**		**Stochastic**		**Deterministic**	
TC	121,756.5	121,763.2	125,374.0	124,988.3	126,475.0	126,178.7
DC	136,988.3	137,023.7	138,427.7	139,942.9	140,798.5	141,276.6
FR	-14,758.5	-14786.4	-12,555.4	-14,477.1	-13,810.9	-14,608.4
E	-473.4	-474.0	-498.2	-477.5	-512.6	-489.5
Horizon	**1-day**					
σ_t	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
Cost Item ($/month)	**Perfect Information**		**Stochastic**		**Deterministic**	
TC	121,796.0	122,234.8	125,570.3	125,825.4	126,322.9	125,977.1
DC	136,983.3	137,537.5	138,427.6	141,113.9	140,611.2	141,220.2
FR	-14,715.6	-14,827.8	-12,371.3	-14,811.7	-13,775.2	-14,761.2
E	-471.7	-474.9	-485.9	-476.8	-513.0	-481.8

Table 2.2: Value of stochastic solution (VSS) and value of perfect information (VPI).

		7-day		**3-day**		**1-day**	
σ_t		**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$	**1**	$\frac{\mathbf{M}}{\mathbf{N}}$
VSS	**($/month)**	1,215.3	1,608.1	1,100.9	1,190.5	752.6	151.7
	(Approx. $/year)	14,583.6	19,297.2	13,210.8	14,286.0	9,031.2	1,820.4
VPI	**($/month)**	4,448.9	4,435.5	4,718.5	4,415.5	4,526.9	3,742.3
	(Approx. $/year)	53,386.8	53,226.0	56,622.0	52,986.0	54,322.8	44,907.6

Table 2.3: Value of the battery under different MPC schemes.

	Without Battery	With Battery			
		Perfect Information MPC	Stochastic MPC	Deterministic MPC	Value of Stochastic MPC
Expected Total Cost ($/month)	138,224	121,794	124,622	126,230	1,608
Expected Savings ($/month)	-	16,430	13,602	11,994	1,608

2.5 Conclusions

In this chapter, we presented a stochastic MPC framework for evaluating the performance of stationary battery systems that provide frequency regulation capacity to the ISO while simultaneously mitigating demand charges from a local utility. Our framework uses a Ledoit-Wolf covariance estimation to generate load realizations from limited historical data. We use real load data from a typical university campus and regulation data from PJM to evaluate the performance of stochastic MPC against deterministic MPC and perfect information MPC. Our simulations demonstrate that stochastic MPC can mitigate large demand charges and better manage frequency regulation commitments. We observe that the length of the prediction horizon and discounting of the demand charge significantly affect economic performance. Notably, we find that stochastic MPC can recover 83% of the value of the battery (which we define as the savings obtained by installing the battery and operating it under perfect information) while deterministic MPC can only recover 73% of such value. Our results also show that important reductions in the payback period of the battery can be achieved by using stochastic MPC.

3

BENCHMARKING STOCHASTIC AND DETERMINISTIC MPC: A CASE STUDY IN STATIONARY BATTERY SYSTEMS

In this chapter, we present a computational framework that integrates forecasting, uncertainty quantification, and model predictive control (MPC) to benchmark the performance of deterministic and stochastic MPC. By means of a battery management case study, we illustrate how off-the-shelf deterministic MPC implementations can suffer significant losses in performance and constraint violations due to their inability to handle disturbances that cannot be adequately represented by mean (most likely) forecasts (Kumar et al., 2019a). We also show that adding constraint back-off terms can help ameliorate these issues but this approach is ad-hoc and does not provide performance guarantees. Stochastic MPC provides a more systematic framework to handle these issues by directly capturing uncertainty descriptions of a wide range of disturbances.

3.1 Introduction

Model predictive control (MPC) revolutionized the field of control due (in large part) to its ability to anticipate and counteract uncertain disturbances in order to maximize performance and satisfy system constraints Qin and Badgwell (2003). In a standard MPC imple-

mentation, a deterministic representation of uncertain disturbances (i.e., a forecast) is used to compute control actions Oldewurtel et al. (2012); Zavala et al. (2009b); Appino et al. (2018). In other words, the forecast acts as a *summarizing statistic* of the entire disturbance uncertainty space. The forecast is typically the most likely realization of the disturbance (usually the mean), and can be conveniently obtained using techniques such as autoregressive models. Deviations of the actual disturbance realization from the forecast will introduce a forecast error, which is counteracted by adjusting the control action at the next time step. This feedback mechanism provides inherent robustness to the MPC controller and seeks to maintain the system within constraints Magni and Scattolini (2007); Limon et al. (2010). This inherent robustness property can be, however, insufficient to avoid constraint violations. For instance, constraint violations might be encountered during the sampling period (when the control action is fixed) as the system faces the actual disturbance realization. Moreover, a large forecast error might put the MPC controller in a state at the next time step in which no feasible control policy exists (i.e., recursive feasibility is lost).

An important observation that we make in this work is that *certain types of disturbances cannot be well-represented using summarizing statistics* and this can trigger inconsistencies in the computation of the MPC control action. This is the case, for instance, when disturbances have zero mean and are multiplicative and/or when they exhibit discrete values (e.g., ON/OFF). To capture such disturbances, it is necessary to use an explicit characterization of the uncertainty space in the MPC formulation. This can be done by using stochastic MPC or robust MPC formulations Kerrigan (2001); Bernardini and Bemporad (2009); Lucia et al. (2014); de la Penad et al. (2005). In this work, we focus our discussion on stochastic MPC formulations because they can capture diverse types of uncertainty characterizations (but many of the benefits discussed can also be obtained with robust MPC such as decision-making for extreme scenarios, however, robust MPC does not cover the probability of occurrences of more likely scenarios of uncertainty).

A wide range of stochastic MPC formulations have been explored in the literature and these vary mostly in: i) the approach used to describe uncertainty, and ii) the mechanism

used to enforce constraints. In terms of i), multi-stage stochastic MPC formulations rigorously handle uncertainty descriptions by capturing the fact that information about the disturbances gets revealed progressively over time and uncertainty evolution is modeled as a scenario tree. Lucia et al. (2013). This enables the computation of an entire control policy but this approach is, in general, intractable due to an exponential explosion in complexity as one discretizes the uncertainty representation (using sampling). An approximate and tractable policy can be obtained by using a so-called two-stage stochastic MPC approach, in which one captures the uncertainty of the disturbances when computing the next immediate control action, but it is assumed that no information is gained about the disturbances in the future. Instead, this approach uses a receding horizon implementation to factor in new information and update the uncertainty description Kumar et al. (2018a). In terms of ii), constraints can be enforced in stochastic MPC by ensuring that they hold under each realization of the disturbances. This approach is computationally flexible but conservative because it is equivalent to enforcing the constraints with probability one. A less conservative approach consists of enforcing the constraints with a lower probability, by using so-called chance constraints Mesbah (2016). Unfortunately, chance-constrained MPC formulations offer less computational flexibility in that they are only tractable in certain specific settings. In this work, we opt for computational flexibility and consider scenario-based, two-stage stochastic MPC formulations. We will see that this approach is sufficient to start realizing significant benefits over deterministic MPC.

In this work, we present a computational framework that integrates disturbance forecasting, uncertainty quantification, and MPC to provide a critical assessment of deterministic and stochastic MPC. Our framework focuses on a case study that arises in the management of stationary battery systems that provide simultaneous frequency regulation (FR) services to the power grid while helping mitigate peak monthly demand charge costs of a university campus. The study uses real disturbance data and a rigorous benchmarking procedure to systematically compare the performance of deterministic and stochastic MPC policies. Our study reveals that FR dispatch signals from the power grid cannot be

properly captured by deterministic MPC formulations and this introduces ineffective control actions that lead to losses in performance and constraint violations (FR commitments cannot be satisfied). This issue can be counteracted using a back-off term that allocates battery reserve capacity but this approach is ad-hoc and prevents full utilization of the battery asset (thus decreasing its value). This indicates that, contrary to common perception, feedback alone is insufficient to counteract certain types of disturbances. Stochastic MPC provides a more systematic framework to account for diverse disturbances, satisfy constraints, and maximize asset value. We derive a quasi-stochastic MPC formulation that uses a coarse representation of the uncertainty space. We show that this simple approach already avoids many of the limitations of deterministic MPC while decreasing the computational complexity of stochastic MPC.

3.2 Computational Framework

In this section, we describe the elements of our computational framework. We begin by providing a description of the decision-making problem faced by the battery. We then describe how data is used to create forecasts and uncertainty characterizations for the disturbances affecting the battery. Then, we describe the MPC formulations under study as well as metrics and procedures used to perform benchmarks. The nomenclature used in this chapter is provided in Section A.1 of Appendix A.

3.2.1 *Preliminaries*

We consider the real-time management of a stationary battery system to determine optimal participation strategies in energy and FR markets operated by an independent system operator (ISO) while simultaneously providing demand-side management for a collection of buildings and with this mitigate monthly demand charges from a utility company (Kumar et al., 2018a). As more intermittent renewable power is injected into the grid, battery systems provide flexibility to the ISOs by providing energy and frequency regulation (FR) ca-

pacity and to the utilities by providing demand-side management capabilities for energy-intensive facilities (e.g. buildings or manufacturing units) (Rastler, 2010; Oudalov et al., 2006). There has been an increased demand for fast dispatchable resources and frequency regulation services by the ISOs in order to mitigate high-frequency renewable fluctuations and stranded power (Kim et al., 2016). Stationary battery systems can also be strategically placed in the network to enhance system-wide performance (He et al., 2015; Sioshansi et al., 2009). The economic benefits provided by stationary battery systems (in the context of FR) have been widely studied in the literature (Mercier et al., 2009; Mohsenian-Rad, 2016; Kottick et al., 1993; Singh et al., 2014). *Simultaneous* FR provision and demand charge mitigation using battery systems has also been studied in Sigrist et al. (2013); White and Zhang (2011); Lucas and Chondrogiannis (2016); Sebastián (2016).

The various cost and revenue components faced by the battery are:

- *Energy Transactions (hourly)*: The battery system purchases energy to recharge and discharges to provide energy for the building and for the FR signal. The energy transactions are charged at the time-varying market energy price.

- *Frequency Regulation Capacity (hourly)*: The ISO compensates the battery for providing an operational band (compensated based on time-varying market FR capacity prices) around a charge/discharge level of the battery (net power of the battery, charged at market energy prices). The ISO can request the battery to dispatch a fraction of the committed capacity based on the grid requirements using a real-time FR signal which is a zero-mean signal with a bounded range of [-1,+1].

- *Demand Charges (monthly)*: The attached load is charged for the peak demand (at a fixed demand charge price) incurred over a month by the utility company.

The combined battery and load system is illustrated in Figure 3.1. This illustrates the interaction between the system, the ISO, and the utility. Section A.1 of Appendix A provides a detailed description of all the variables and quantities used in the battery management framework.

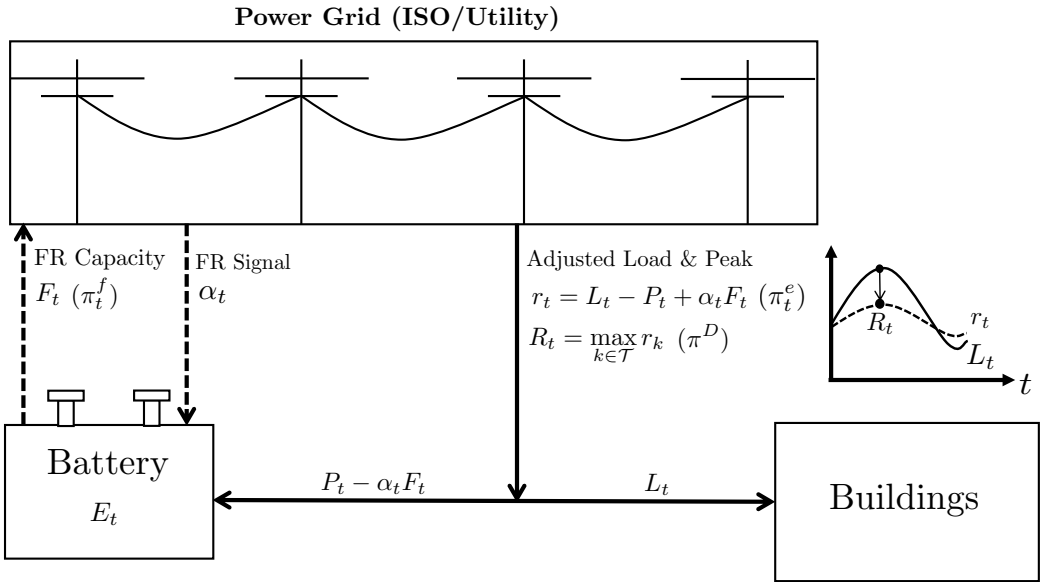

Figure 3.1: Combined battery and attached load system.

The proposed MPC implementations solve discrete-time optimal control problems at every time $t \in \mathbb{Z}_+$ with a prediction horizon $\mathcal{T} := \{t+1, t+2, \ldots, t+N\}$. Here, N is the length of the prediction horizon and index t corresponds to the time instant $t \cdot h$, where $h = 1$ hour. Quantities with subindex t are *held constant* over the time interval $[(t-1) \cdot h, t \cdot h]$. The MPC schemes are run in a receding-horizon manner (updated every hour) over a horizon denoted by $\mathcal{M} := \{0, \ldots, M\}$ (assumed here to be one month).

The FR capacity F_t represents a *symmetric band* around the middle point P_t, which is the battery net charge/discharge rate. The notation $P_{\mathcal{T}}$ is used to denote the net rate over the prediction horizon $\mathcal{T}$. A similar notation is used to denote all time trajectories. The battery commits a net discharge level P_t and an FR band of length $2F_t$ to ISO over the corresponding period. The actual power requested by the ISO is $\alpha_t F_t$, where $\alpha_t \in [-1, 1]$ is an exogenous FR signal coming from the ISO. A value of $\alpha_t = 1$ indicates that the ISO requests all the FR capacity by charging the battery while $\alpha_t = -1$ indicates that the ISO requests all the capacity by discharging. The FR capacity F_t provided by the battery is remunerated at price π_t^f (this flexibility is a valuable asset for the ISO). The net rate P_t is also used to shape the residual load $r_t = L_t - P_t + \alpha_t F_t$, which is the load that is paid for to

the utility company. Here, L_t is the exogenous attached load coming from the buildings. The residual load incurs an hourly cost at price π_t^e as well as a monthly peak cost, which is charged based on the maximum residual load $\max_{t\in\mathcal{M}} r_t$ at price π^D.

We thus see that a battery management system must *strategically* decide how to use the battery storage capacity to allocate P_t and F_t in order to minimize energy costs and demand charges, maximize FR revenue, and to maintain suitable levels for stored energy.

3.2.2 *Forecasting and Uncertainty Quantification*

The decision to allocate P_t and F_t is complicated by the fact that the energy and FR prices $\pi_t = (\pi_t^e, \pi_t^f)$, the attached load L_t, and the FR signal α_t are all highly uncertain. These disturbances are modeled as random variables with realizations denoted by index ξ. At time t, a realization of the uncertain disturbances over a forecast horizon $\mathcal{T}$ is denoted as $L_\mathcal{T}(\xi), \pi_\mathcal{T}(\xi)$ and $\alpha_\mathcal{T}(\xi)$. The forecast trajectory is a specific realization (usually the one with highest probability) and is denoted as $\hat{L}_\mathcal{T}, \hat{\alpha}_\mathcal{T}$, and $\hat{\pi}_\mathcal{T}$. To enable a compact representation, we encapsulate all disturbances in the random vector $d_\mathcal{T} := (\pi_\mathcal{T}, \alpha_\mathcal{T}, L_\mathcal{T})$. We denote a realization of the disturbance $d_\mathcal{T}$ vector as $d_\mathcal{T}(\xi)$ and we denote the entire set of disturbance realizations using the notation $d_\mathcal{T}(\Xi)$.

We assume that the load and energy prices are Gaussian (normal) variables of the form $L_\mathcal{T} \sim \mathcal{N}(\hat{L_\mathcal{T}}, \Sigma_\mathcal{T}^L)$ and $\pi_\mathcal{T}^e \sim \mathcal{N}(\hat{\pi_\mathcal{T}^e}, \Sigma_\mathcal{T}^e)$, respectively. Time correlations in these disturbances are implicitly captured in the covariance matrices. The FR signal α_t is assumed to be Gaussian and of the form $\alpha_t \sim \mathcal{N}(\hat{\alpha}_t, \sigma_\alpha^2)$. The FR price π_t^f is assumed to be log-normal with $\log(\pi_t^f) \sim \mathcal{N}(\log(\hat{\pi}_t^f), \sigma_{\pi^f}^2)$. Note that this implicitly assumes that no time correlations are present in the FR signal and FR prices. A justification for the choice of random variables to represent the data is provided in recent work (Kumar et al., 2018a). The mean and variances for the uncertainty descriptions are updated using a receding-horizon scheme that uses historical and real-time data (as it becomes available). We denote the observed (actual) disturbance history at time t as $d_\mathcal{H}$, where $\mathcal{H} := \{t-H, t-H+1, \ldots, t\}$.

In this work, we assume that these random variables are independent. This drastically simplifies forecasting and uncertainty quantification procedures. The use and impact of different uncertainty characterizations is a non-trivial but interesting topic of future work.

The mean and covariance matrices for $L_{\mathcal{T}}$ and $\pi^e_{\mathcal{T}}$ are obtained using autoregressive (AR) models. Specifically, we use time series models of the form:

$$L_t = \sum_{k=1}^{H} \phi^L_k L_{t-k} + c^L + \epsilon^L_t \tag{3.2.1a}$$

$$\pi^e_t = \sum_{k=1}^{H} \phi^e_k \pi^e_{t-k} + c^e + \epsilon^e_t, \tag{3.2.1b}$$

where H is the order of the model, ϕ^L_t, ϕ^T_t, c^L and c^e are coefficients that are learned (estimated) from the historical data, and ϵ^L_t and ϵ^e_t are noise sequences. The mean (most likely) forecasts $\hat{L}_{\mathcal{T}}, \hat{\pi}^e_{\mathcal{T}}$ are obtained by using the maximum likelihood estimates of the coefficients. Maximum likelihood procedures also provide the covariance matrices $\Sigma^L_{\mathcal{T}}$ and $\Sigma^e_{\mathcal{T}}$. Explicit techniques for doing this are provided in Box et al. (2015). In this work, we use standard procedures provided in the `R` software package. The disturbance history $d_{\mathcal{H}}$ is also used to compute sample estimates for $\hat{\alpha}$, σ_α, $\hat{\pi}^f$, and $\sigma^2_{\pi^f}$. With this information, we generate a set of disturbance realizations $\bar{\Xi}$ by sampling from the corresponding probability density functions. This set of realizations used for the stochastic MPC formulation and the mean forecast $\hat{d}_{\mathcal{T}} = \mathbb{E}[d_{\mathcal{T}}(\bar{\Xi})]$ is used for the deterministic MPC formulation.

3.2.3 *Deterministic MPC*

The deterministic MPC controller uses the mean forecast $\hat{d}_{\mathcal{T}}$ to find the control policy $P_{\mathcal{T}}, F_{\mathcal{T}}$ that maximizes the total *forecast* net profit (minimize the total net cost) over the prediction horizon $\mathcal{T}$. This is done by solving the optimization problem:

$$\min \sum_{k\in\mathcal{T}} \hat{\pi}^e_k r_k - \sum_{k\in\mathcal{T}} \hat{\pi}^f_k F_k + \pi^D \max_{k\in\mathcal{T}} r_k. \tag{3.2.2a}$$

$$\text{s.t. } P_k + F_k \leq \overline{P},\ k \in \mathcal{T} \tag{3.2.2b}$$

$$P_k - F_k \geq -\underline{P},\ k \in \mathcal{T} \tag{3.2.2c}$$

$$E_k = E_{k-1} - P_k + \hat{\alpha}_k F_k,\ k \in \mathcal{T} \tag{3.2.2d}$$

$$\rho F_k \leq E_{k-1} \leq \overline{E} - \rho F_k,\ k \in \mathcal{T} \tag{3.2.2e}$$

$$\rho F_k \leq E_k \leq \overline{E} - \rho F_k,\ k \in \mathcal{T} \tag{3.2.2f}$$

$$-\overline{\Delta P} \leq P_k - P_{k-1} \leq \overline{\Delta P},\ k \in \mathcal{T} \tag{3.2.2g}$$

$$r_k = \hat{L}_k - P_k + \hat{\alpha}_k F_k,\ k \in \mathcal{T} \tag{3.2.2h}$$

$$\max_{k \in \mathcal{T}} r_k \geq R_t \tag{3.2.2i}$$

$$P_k + F_k \leq \hat{L}_k,\ k \in \mathcal{T} \tag{3.2.2j}$$

$$0 \leq E_k \leq \overline{E},\ k \in \mathcal{T} \tag{3.2.2k}$$

$$-\underline{P} \leq P_k \leq \overline{P},\ k \in \mathcal{T} \tag{3.2.2l}$$

$$0 \leq F_k \leq \overline{P},\ k \in \mathcal{T} \tag{3.2.2m}$$

As in a standard receding-horizon scheme, only the first element of the control policy (P_{t+1}, F_{t+1}) is implemented in the system. The actual disturbance realized $d_{t+1}(\xi)$ will tend to deviate from the forecast $\hat{d}_{t+1}$. As a result, the state at time $t+1$ will differ from that predicted by the MPC controller. In particular, the state of charge (SOC) of the battery evolves as $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$ and the residual load as $r_{t+1}(\xi) = L_{t+1}(\xi) - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$. To account for the error in the prediction, the horizon is shifted forward to update the disturbance forecast using the new data history. In this work, we assume that the disturbances can be measured exactly at each time step and that these in turn can be used to exactly update the current state. Therefore, no state estimation procedure is needed (only a history of disturbance data is required in order to develop prediction models for the disturbance).

The parameter R_t is the carryover peak demand (largest demand seen so far in the month). This highlights that the peak residual demand acts as a state of the system. In particular, the carryover peak demand is given by $R_{t+1} = \max\{R_t, r_{t+1}(\xi)\}$. For simplicity, we

define the control action generated from the solution of the problem as $u_{t+1} = (P_{t+1}, F_{t+1})$ and the state as $x_{t+1} = (E_{t+1}, R_{t+1})$.

Constraints (3.2.2b)-(3.2.2c) are net discharge/charge capacity limits for the battery. Constraints (3.2.2e)-(3.2.2f) play a key role in this formulation as they impose a reserve (buffer) for the battery storage to account for variations in the FR signal. The parameter $\rho \in \mathbb{R}_+$ determines the buffer size: when $\rho = 0$, the MPC formulation is a standard formulation that aims to exploit the entire feasible region to minimize cost while a larger value of ρ allocates more reserves and thus limits the utilization of the battery for cost minimization. Constraint (3.2.2g) imposes limits on subsequent net ramping rates to protect the battery lifetime. Constraints (3.2.2h) and (3.2.2i) are used to compute the peak residual demand. Constraint (3.2.2j) prevents the battery from selling back electricity to the grid (the battery can only purchase). The rest of the constraints are capacity limits.

From the structure of the optimization formulation, we highlight that the mean FR signal always appears as a product with the actual FR capacity. As a result, when a mean forecast of this signal is used (which is close to zero), the FR capacity $F_{\mathcal{T}}$ has a small influence on the constraints and thus the MPC controller tends to push this capacity to a large value in order to maximize profit. We will see that this effect results in overcommitment of FR capacity that the controller is unable to meet, resulting in FR shortfall and constraint violations.

We denote the optimization problem solved in deterministic MPC at time t as $\mathcal{P}^{det}(\hat{d}_{\mathcal{T}}, x_t)$, where the arguments are the current state, control information, and disturbance forecast needed to solve the problem.

We summarize the receding-horizon scheme for the deterministic MPC run over a simulation horizon $\mathcal{M}$:

1. START at $t = 0$ with E_0 and $R_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:
2. Use disturbance history $d_{\mathcal{H}}$ to obtain forecast $\hat{d}_{\mathcal{T}}$.
3. Solve $\mathcal{P}^{det}(x_t, \hat{d}_{\mathcal{T}})$ to obtain commitments $u_{t+1} = (P_{t+1}, F_{t+1})$.

4. Implement the decisions, u_{t+1} over $(t, t+1)$.

5. Update the states using the actual *realized* disturbances $d_{t+1}(\xi)$ as $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$ and $R_{t+1} = \max\{R_t, r_{t+1}(\xi)\}$ with $r_{t+1}(\xi) = L_{t+1} - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$.

6. Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN to Step 2.

In this scheme, the actual realized disturbances are obtained from a set of validation scenarios $\tilde{\Xi}$. These scenarios are generated from actual disturbance data (not from the forecast). The performance of the MPC controller thus depends on the selection of the realized disturbances and is thus random. We thus run the scheme for the entire set of validation samples to obtain probability distributions for diverse performance metrics.

3.2.4 *Stochastic MPC*

In stochastic MPC, uncertainty representations for disturbances $D_{\mathcal{T}}$ are captured in the optimization formulation by using multiple realizations $d_{\mathcal{T}}(\xi)$, $\xi \in \bar{\Xi}$. The net power P_{t+1} and FR commitments F_{t+1} are defined over the next immediate trading hour as *here-and-now decisions* that need to be made prior to observing uncertainty. The battery power and FR capacity trajectories $P_k(\xi)$ and $F_k(\xi)$ at subsequent times $k \in \mathcal{T} \setminus \{t+1\}$ are modeled as *recourse decisions* that can be corrected according to the disturbance realization. The residual and peak demands are also recourse decisions that are expressed as $r_{\mathcal{T}}(\xi)$ and $\max_{k\in\mathcal{T}} r_k(\xi)$. The SOC at time $t+1$ only depends on the previous storage E_t and on P_{t+1}. Consequently, E_{t+1} is also a here-and-now variable. The rest of the trajectory $E_k(\xi)$ for $k \in \mathcal{T} \setminus \{t+1\}$ is recourse variable because $P_k(\xi)$ are recourse variables for $k \in \mathcal{T} \setminus \{t+1\}$.

We use $\mathcal{P}^{sto}(d_{\mathcal{T}}(\bar{\Xi}), x_t)$ to denote the optimization problem solved in stochastic MPC at time t. The variables and constraints of the stochastic MPC formulation are the same as those of the deterministic counterpart but are replicated for the set of realizations $\xi \in \bar{\Xi}$. We use the non-anticipativity constraints to enforce the fact that the control actions P_{t+1} and F_{t+1} are here-and-now (commitment) decisions that need to be implemented in the

system.

$$\min \mathbb{E}\left[\sum_{k\in\mathcal{T}} \pi_k^e(\xi) r_k(\xi) - \sum_{k\in\mathcal{T}} \pi_k^f(\xi) F_k(\xi) + \pi^D \max_{k\in\mathcal{T}} r_k(\xi)\right] \tag{3.2.3a}$$

$$\text{s.t. } P_k(\xi) + F_k(\xi) \leq \overline{P},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3b}$$

$$P_k(\xi) - F_k(\xi) \geq -\underline{P},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3c}$$

$$E_k(\xi) = E_{k-1}(\xi) - P_k(\xi) + \alpha_k(\xi) F_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3d}$$

$$\rho F_k(\xi) \leq E_{k-1}(\xi) \leq \overline{E} - \rho F_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3e}$$

$$\rho F_k(\xi) \leq E_k(\xi) \leq \overline{E} - \rho F_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3f}$$

$$-\overline{\Delta P} \leq P_k(\xi) - P_{k-1}(\xi) \leq \overline{\Delta P},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3g}$$

$$r_k(\xi) = L_k(\xi) - P_k(\xi) + \alpha_k(\xi) F_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3h}$$

$$P_k(\xi) + F_k(\xi) \leq L_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3i}$$

$$\max_{k\in\mathcal{T}} r_k(\xi) \geq R_t,\ \xi \in \bar{\Xi} \tag{3.2.3j}$$

$$0 \leq E_k(\xi) \leq \overline{E},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3k}$$

$$-\underline{P} \leq P_k(\xi) \leq \overline{P},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3l}$$

$$0 \leq F_k(\xi) \leq \overline{P},\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{3.2.3m}$$

$$P_{t+1}(\xi) = P_{t+1}(\xi'),\ \xi \neq \xi', \xi, \xi' \in \bar{\Xi} \tag{3.2.3n}$$

$$F_{t+1}(\xi) = F_{t+1}(\xi'),\ \xi \neq \xi', \xi, \xi' \in \bar{\Xi} \tag{3.2.3o}$$

In the above formulation, the expected value $\mathbb{E}[\cdot]$ is defined over the set of scenarios $\bar{\Xi}$. From the structure of the model, we see that we use multiple realizations of the FR signal $\alpha_{\mathcal{T}}(\bar{\Xi})$. As a result, the stochastic MPC controller is expected to be better able to meet FR commitments and constraints.

We can now summarize the stochastic MPC scheme run over the simulation horizon $\mathcal{M}$:

1. START at $t = 0$ with E_0 and $R_0 = 0$ given. REPEAT for $t \in \mathcal{M}$:

2. Use disturbance history $d_{\mathcal{H}}$ to obtain forecast realizations $d_{\mathcal{T}}(\bar{\Xi})$.

3. Solve $\mathcal{P}^{sto}(x_t, d_{\mathcal{T}}(\bar{\Xi}))$ to obtain commitments $u_{t+1} = (P_{t+1}, F_{t+1})$.

4. Implement the decisions, u_{t+1} over $(t, t+1)$.

5. Update the states using the actual *realized* disturbances $d_{t+1}(\xi)$ as $E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$ and $R_{t+1} = \max\{R_t, r_{t+1}(\xi)\}$ with $d_{t+1}(\xi) = L_{t+1} - P_{t+1} + \alpha_{t+1}(\xi)F_{t+1}$.

6. Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN to Step 2.

In this scheme, the actual realized disturbances are also obtained from the set of validation samples $\tilde{\Xi}$. Importantly, these validation samples differ from the realizations used in the MPC controller formulation $\hat{\Xi}$. By running the stochastic MPC scheme for all validation samples, we can compute probability distributions for performance metrics that are compared with those from deterministic MPC. This systematic procedure ensures fair comparisons between different MPC implementations.

3.2.5 *Perfect Information MPC*

To evaluate the *ideal* performance of MPC, we also consider a *perfect information MPC* implementation. Under perfect information MPC, we compute commitment policies $P_t(\xi)$ and $F_t(\xi)$ at every time t for discharge/charge power and FR capacity, respectively for *every realization* $\xi \in \tilde{\Xi}$ of the loads, prices, and FR signals. These policies can be computed by removing the nonanticipativity constraints (3.2.3n) and (3.2.3o) from the stochastic MPC formulation (3.2.3) and replacing the disturbance forecast model with the true disturbance signals corresponding to each realization, and implementing the same scheme as the stochastic MPC run over a month period $\mathcal{T}_M$.

3.2.6 *Handling Constraint Violations*

Poor forecasts can lead MPC controllers to violate the constraints during the transition $(t, t+1)$. To capture this issue, an auxiliary MPC controller that uses battery reserves to

make power corrections is used. At time t, this controller solves the feasibility restoration problem:

$$\min_{\Delta P} |\Delta P| \tag{3.2.4a}$$

$$\text{s.t. } P_{t+1} + F_{t+1} + \Delta P \leq \overline{P}, \tag{3.2.4b}$$

$$P_{t+1} - F_{t+1} + \Delta P \geq -\underline{P}, \tag{3.2.4c}$$

$$E_{t+1} = E_t - P_{t+1} + \alpha_{t+1}(\xi) F_{t+1} + \Delta P \tag{3.2.4d}$$

$$-\overline{\Delta P} \leq (P_{t+1} - \Delta P) - P_t \leq \overline{\Delta P}, \tag{3.2.4e}$$

$$P_{t+1} + F_{t+1} + \Delta P \leq L_{t+1}(\xi), \tag{3.2.4f}$$

$$0 \leq E_{t+1} \leq \overline{E}, \tag{3.2.4g}$$

This formulation uses the realized values of the FR signal and the load over the time interval $(t, t+1)$. The feasibility restoration problem seeks to find a net rate correction ΔP that satisfies the constraints. If the auxiliary controller fails to remain feasible even after solving the restoration problem, we assume that no action is taken (no power discharged and no FR commitment is made). In other words, we set $P_{t+1} = F_{t+1} = 0$ and correct the states E_{t+1} and R_{t+1} accordingly. This leads to FR shortfall, because the controller fails to deliver the committed capacity F_{t+1}.

3.2.7 *Benchmarking Procedure*

To distinguish the policies obtained from the three MPC schemes, we denote the policies obtained from the deterministic MPC as P_t^{det}, F_t^{det} and $r_t^{det}(\xi)$, those from stochastic MPC scheme as P_t^{sto}, F_t^{sto} and $r_t^{sto}(\xi)$, and those from perfect information MPC as $P_t^{perf}(\xi)$, $F_t^{perf}(\xi)$ and $r_t^{perf}(\xi)$ (corresponding to each realization ξ).

Each realization in the validation set $\tilde{\Xi}$ generates a monthly cost for the MPC controllers. The monthly cost under stochastic MPC for a given realization $\xi \in \tilde{\Xi}$ and under a

given closed-loop policy $F_{\mathcal{M}}, P_{\mathcal{M}}, r_{\mathcal{M}}$ is given by:

$$\Phi^{det}(\xi) := \sum_{t\in\mathcal{M}} \pi_t^e(\xi)(\alpha_t(\xi)F_t^{det} - P_t^{det}) - \sum_{t\in\mathcal{M}} \pi_t^f(\xi)F_t^{det} + \pi^D \max_{t\in\mathcal{M}} r_t^{det}(\xi). \tag{3.2.5}$$

The monthly cost for deterministic MPC is denoted as $\Phi^{det}(\xi)$ and is defined as in (3.2.5). Similarly, the monthly cost for stochastic MPC is denoted as $\Phi^{sto}(\xi)$ and defined in the same way as (4.2.7) but with its corresponding closed-loop policy P_t^{sto}, F_t^{sto}, and $r_t^{sto}(\xi)$. The monthly cost for *perfect information MPC* is:

$$\Phi^{perf}(\xi) := \sum_{t\in\mathcal{M}} \pi_t^e(\xi)(\alpha_t(\xi)F_t^{perf}(\xi) - P_t^{perf}(\xi)) - \sum_{t\in\mathcal{M}} \pi_t^f(\xi)F_t^{perf}(\xi) + \pi^D \max_{t\in\mathcal{M}} r_t^{perf}(\xi). \tag{3.2.6}$$

The costs for the different validation realizations are used to create empirical probability distributions and cumulative probability distributions and to compute statistics such as expected costs $\mathbb{E}[\Phi^{det}(\tilde{\Xi})]$, $\mathbb{E}[\Phi^{sto}(\tilde{\Xi})]$, and $\mathbb{E}[\Phi^{perf}(\tilde{\Xi})]$.

Of particular interest in our context is a metric that we call the *expected value of the battery*. To compute this value, we evaluate the total cost under the assumption that no battery has been installed (and thus FR profit and peak shaving capabilities are not available). The cost for particular validation realization $\xi \in \tilde{\Xi}$ is denoted as $\Phi^{nobat}(\xi)$. The ideal expected value of the battery (the ideal expected savings) is defined as $\text{VB}^{perf}(\xi) := \Phi^{nobat}(\xi) - \Phi^{perf}(\xi)$. The value of the battery under stochastic MPC is $\text{VB}^{sto}(\xi) := \Phi^{nobat}(\xi) - \Phi^{sto}(\xi)$ and under deterministic MPC is $\text{VB}^{det}(\xi) := \Phi^{nobat}(\xi) - \Phi^{det}(\xi)$. As in the case of cost, the realizations are used to obtain probability distributions and to compute statistics such as $\mathbb{E}[\text{VB}^{perf}(\tilde{\Xi})]$, $\mathbb{E}[\text{VB}^{sto}(\tilde{\Xi})]$, and $\mathbb{E}[\text{VB}^{det}(\tilde{\Xi})]$. The value of the battery is a metric that reflects loss/gains in asset value due to the use of better control policies. We also consider the value of the stochastic MPC, which is defined as $\text{VSMPC}(\xi) := \text{VB}^{sto}(\xi) - \text{VB}^{det}(\xi)$, and the expected value of stochastic MPC as $\mathbb{E}[\text{VSMPC}(\tilde{\Xi})]$.

Figure 3.2 provides a sketch of the computational framework. The control decisions

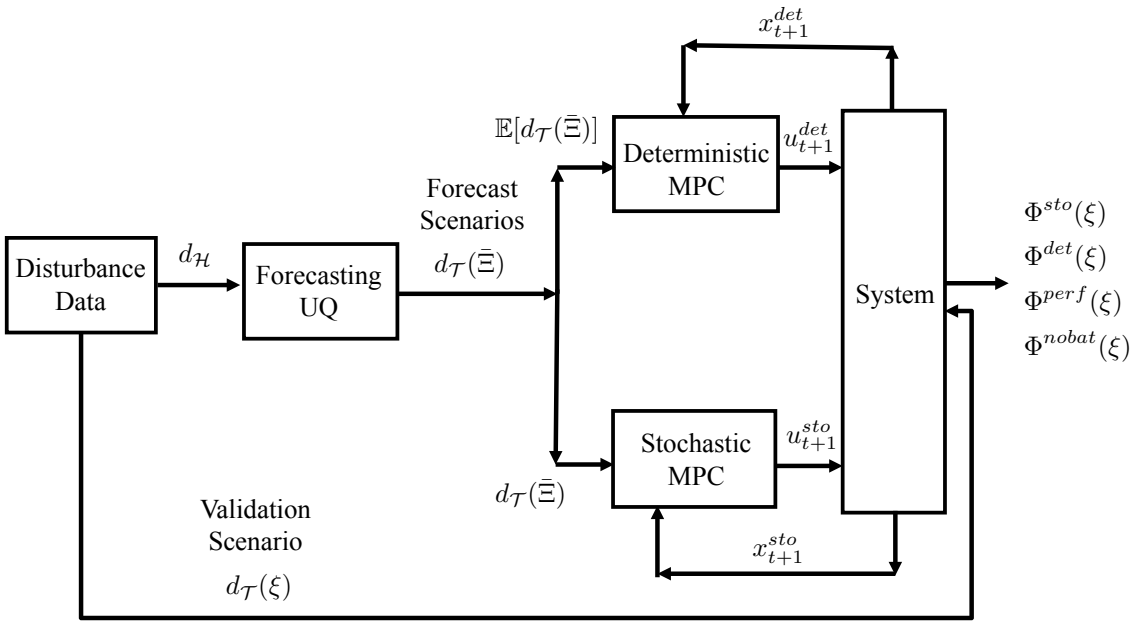

Figure 3.2: Sketch of computational framework.

made by the deterministic MPC at time t are denoted by u_t^{det} while those made by the stochastic MPC at time t are denoted by u_t^{sto}. The state of the system is denoted by x_t.

3.3 Assessment Results

We now present closed-loop simulation results for deterministic, stochastic, and perfect information MPC for an entire month of operation. The controllers use a prediction horizon of 168 hours (seven days), which is chosen based on the observation that the data for the load and electricity prices exhibit weekly periodicity (e.g., high load on weekdays and low load on weekends). In other words, a horizon of seven days captures periodic effects (Kumar et al., 2018a). The stochastic MPC problem contained 100 forecast realizations and a total of 25,300 variables and 92,600 constraints. The realizations are obtained using Monte Carlo sampling. Other sampling techniques such as sparse grids and Latin hypercube sampling can also be accommodated in the proposed stochastic MPC framework to reduce the number of samples but these approaches do not scale to high-dimensional uncertainty

spaces as those considered in time-dependent applications (Renteria et al., 2018). A total of 200 closed-loop monthly runs were performed for each MPC implementation (using the same validation scenarios). The number of forecast and validation scenarios are chosen considering a trade-off between performance and computational cost. The simulations were run on a 32-core machine with Ubuntu 14.04, Intel Xeon 2.30 GHz processors, and 188 GB RAM. The schemes are implemented in `Julia` and leverage the algebraic modeling capabilities of `JuMP` (Dunning et al., 2017), and the optimization problems are solved in extensive form using `Gurobi 8.0` (no decomposition schemes were implemented). Each MPC instance takes about 1 second for deterministic MPC and about 3-4 seconds for stochastic MPC to solve. The monthly closed-loop simulation required approximately 10 minutes for deterministic MPC and 35 minutes for stochastic MPC. The computational times comprise forecasting, optimization solution, and feasibility checks. These times can be made independent of the number of validation scenarios by using parallelization. Time-series disturbance forecasts were obtained using autoregressive models of order 168. The statistics software package `R` is used to obtain such forecasts.

3.3.1 *Closed-Loop Behavior*

We compare the closed-loop policies for both deterministic and stochastic MPC. In these results, an FR buffer of $\rho = 0$ is used for stochastic MPC and a buffer of 10% ($\rho = 0.1$) is used for deterministic MPC. Figure 3.3 provides a snapshot for a given validation scenario for deterministic MPC. Here, we note that the controller uses the buffer fairly frequently to counteract the FR signal. This is because, as can be seen, the mean FR signal forecast is not an appropriate representation of the actual realized signals. This is in sharp contrast with the mean forecasts for the loads and electricity prices, which can be forecast quite accurately. Figure 3.4 shows a snapshot for stochastic MPC under a given validation scenario. Here, we note that the FR signal realizations capture the variability well and thus the controller does not require an explicit FR buffer (which results in better utilization of

the battery). An animation of the closed-loop performance of deterministic and stochastic MPC can be found at https://goo.gl/QAHXpB.

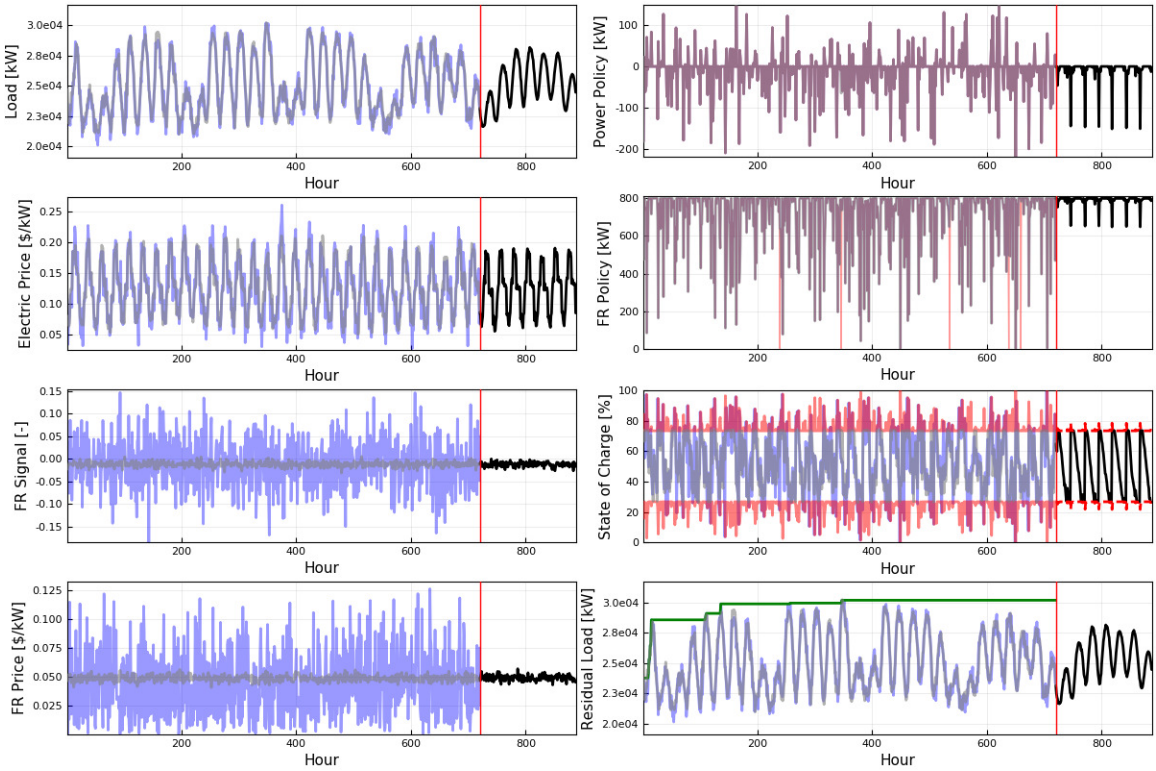

Figure 3.3: Closed loop profile for deterministic MPC with $\rho = 0.1$. Black lines represent forecasts and model predictions. Blue lines represent actual realizations. For the control policies, red lines represent the actual implemented policy. For the state of charge, red lines represent the FR buffer. For the residual load, the green line represents the running peak R_t.

Figure 3.5 compares the cumulative FR shortfall that results from infeasibilities. Clearly, even with a buffer, deterministic MPC incurs more shortfall. FR shortfall arises in stochastic MPC because the controller uses only a finite number of realizations that do not necessarily match the validation realizations. We can see, however, that the controller is significantly more robust than deterministic MPC (the total shortfall is reduced by over 200%).

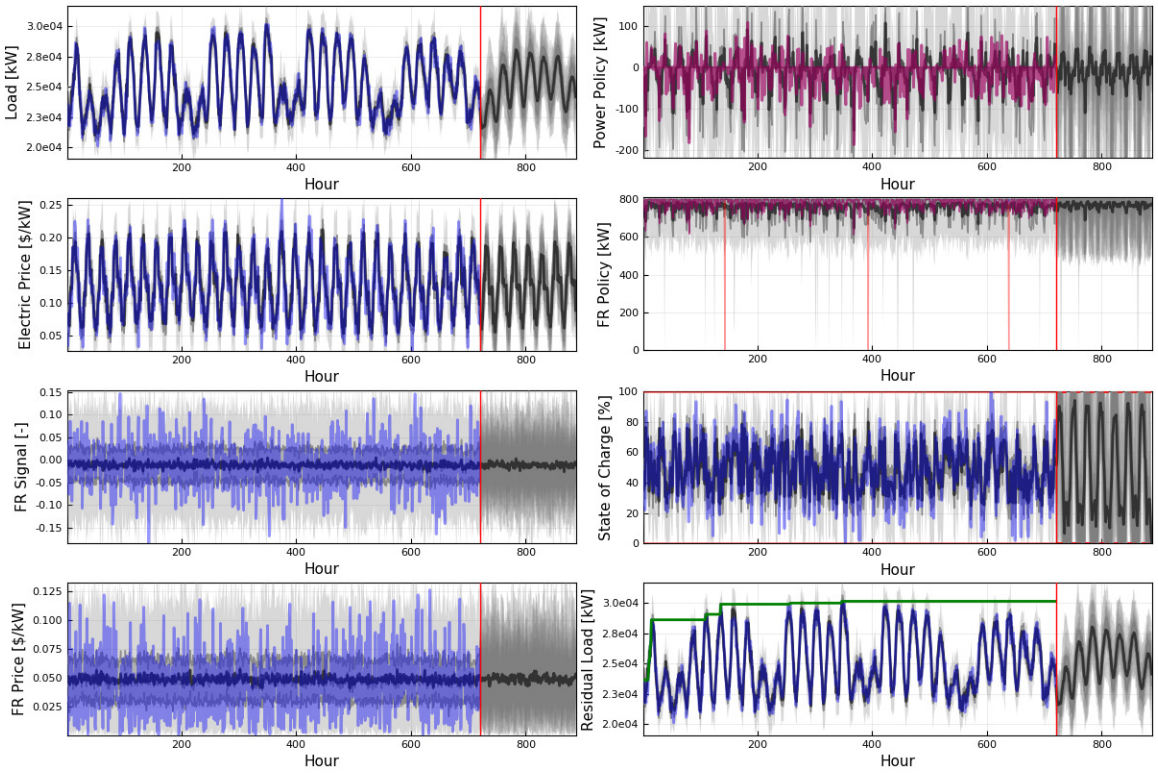

Figure 3.4: Closed loop profile for stochastic MPC with $\rho = 0$. The grey regions represent uncertainty forecasts with black representing mean forecasts. Blue lines represent realized observations or committed policies. For the control policies, red lines represent the actual implemented policy and usually overlap with committed policy. For the residual load, the green line represents the peak observed residual.

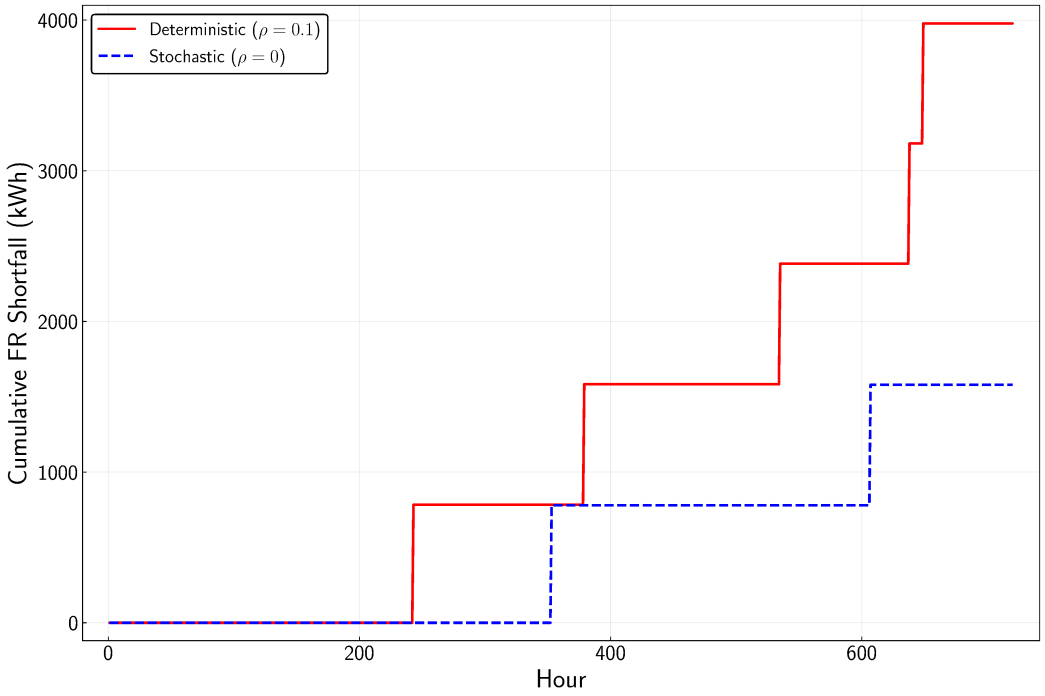

Figure 3.5: Comparison of FR shortfall for deterministic and stochastic MPC for a validation scenario.

3.3.2 *Economic Performance*

The probability and cumulative distributions of total costs, peak demands, and the value of the battery for the 200 validation scenarios are shown in Figures 3.6-3.8. From the distributions of the demand charges, we see that deterministic and stochastic MPC achieve very similar values. This indicates that the benefits of stochastic MPC do not result from demand charge mitigation but from better utilization of the battery to collect FR revenue. This can also be seen in Table 3.1, where we present the expected cost, battery value, and FR shortfall comparisons. The expected cost of the system without using the battery is 136,435 \$/month, which shows the large energy costs associated with this system. We also see that the expected value of the battery of deterministic MPC is improved by 8% by using stochastic MPC. This represents half of the total possible improvement (obtained with perfect information MPC).

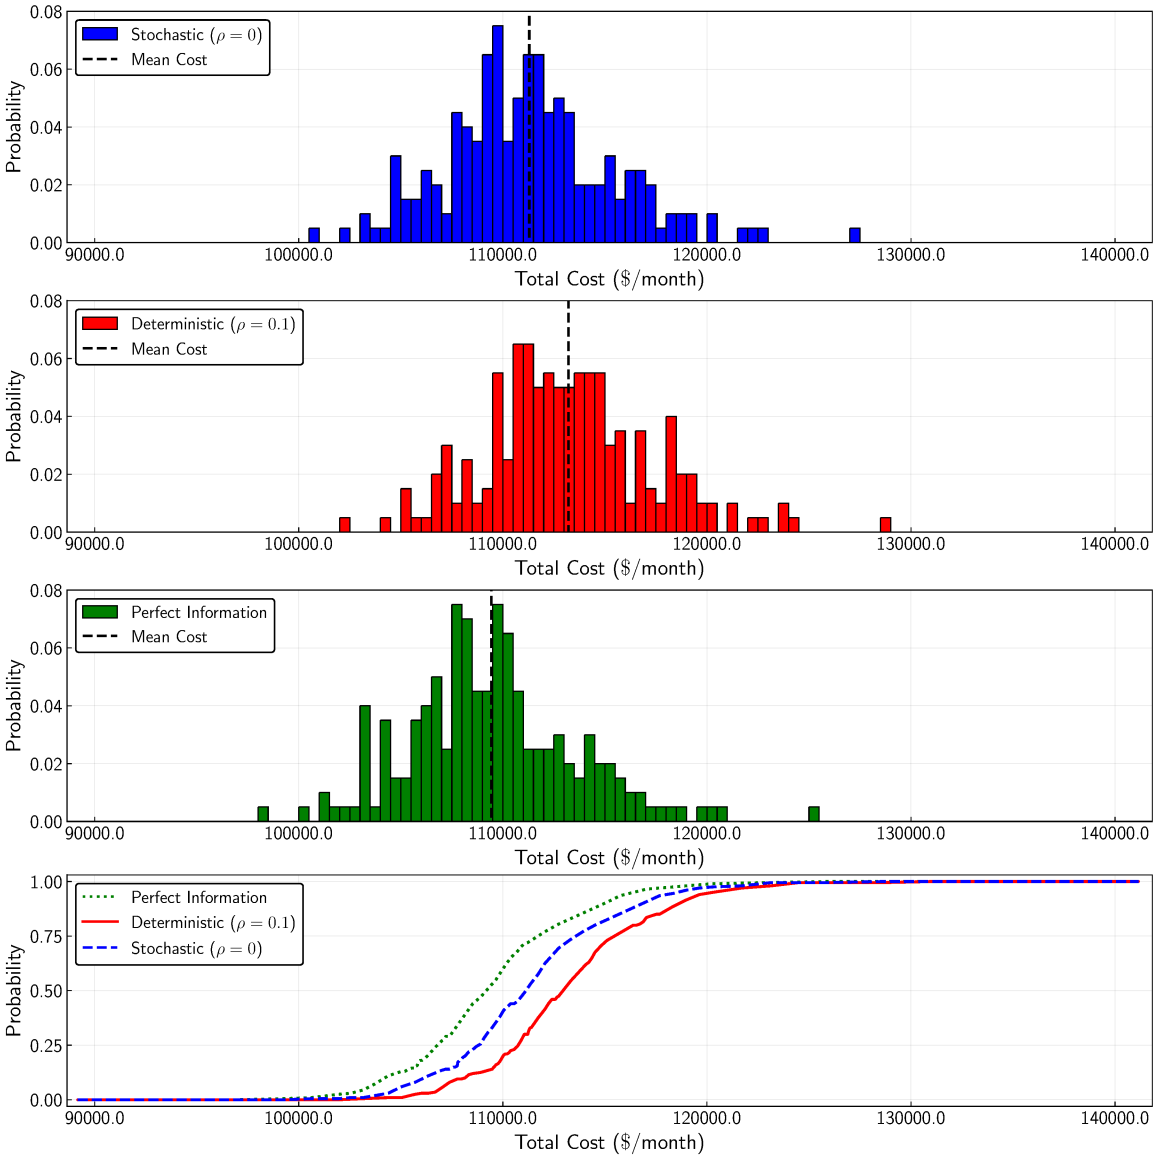

Figure 3.6: Probability and cumulative distributions for total costs.

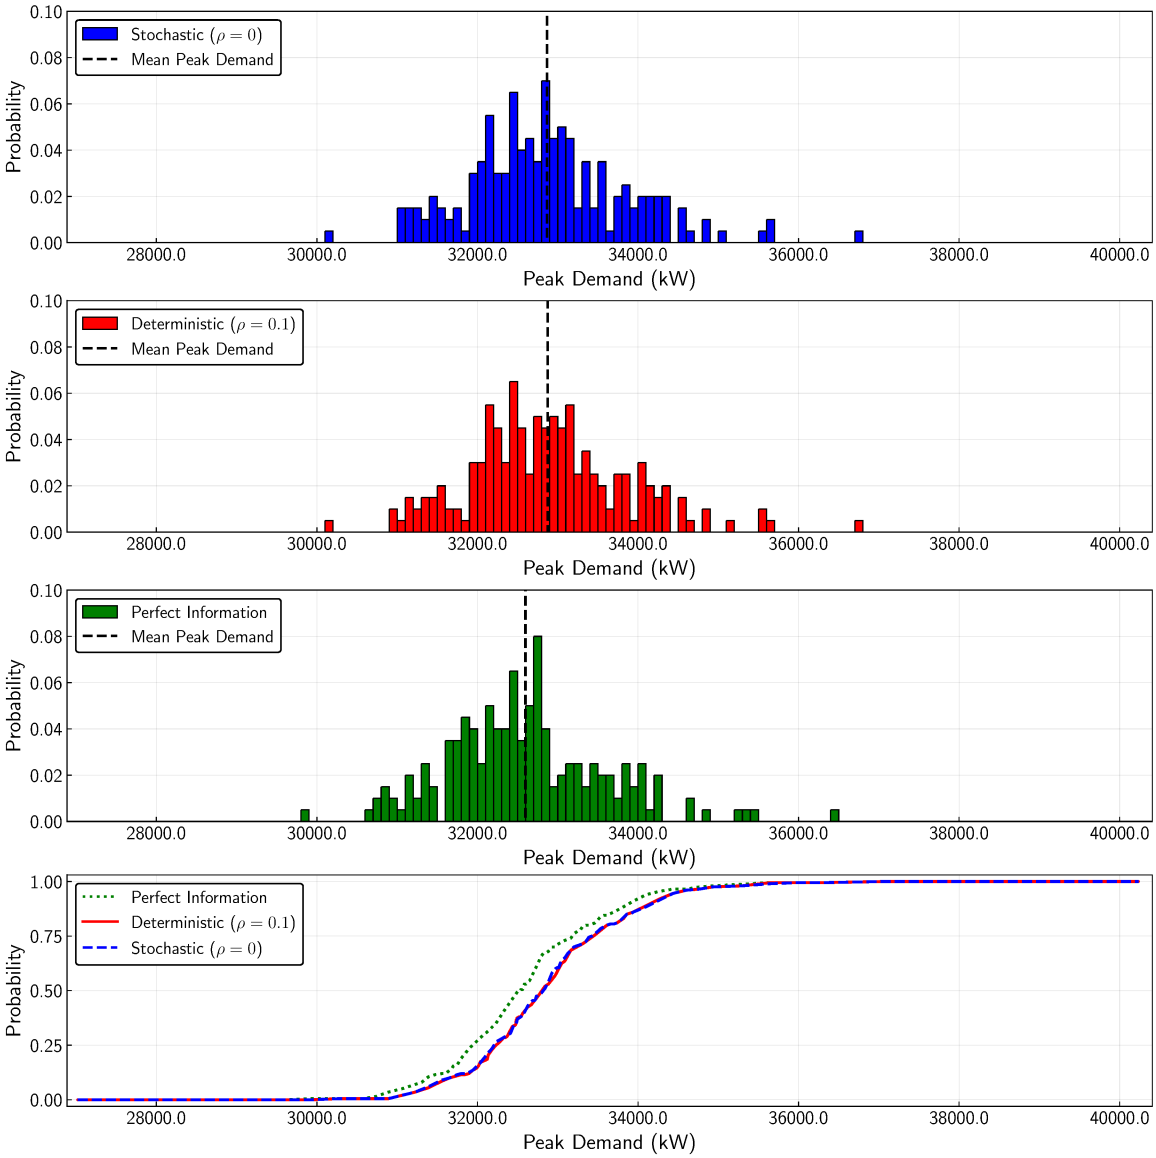

Figure 3.7: Probability and cumulative distributions for demand charges.

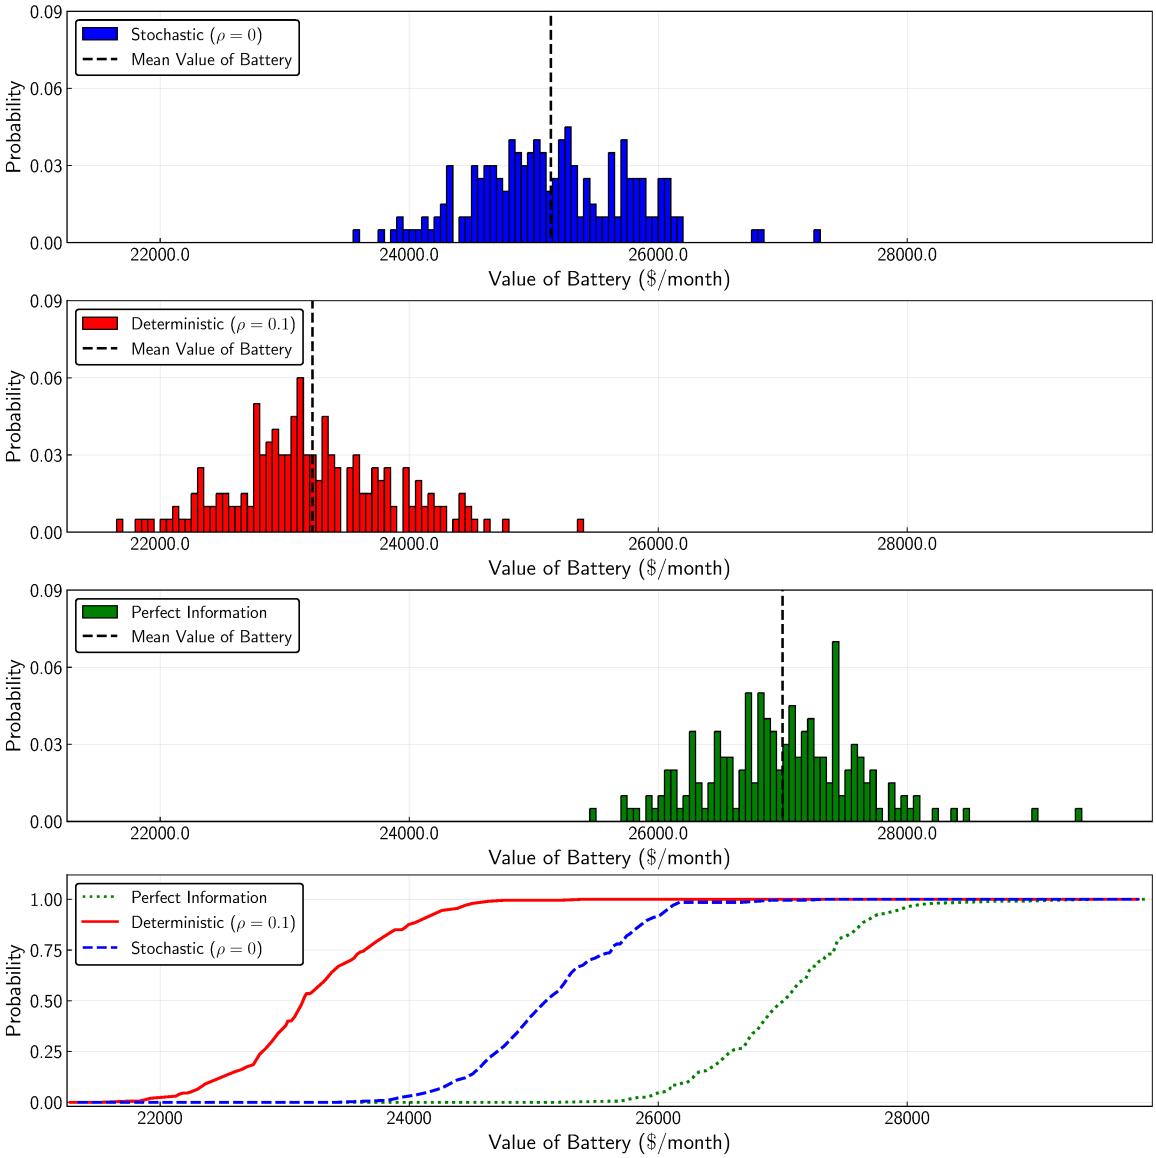

Figure 3.8: Probability and cumulative distributions for the value of the battery.

Table 3.1: Economic performance analysis for different MPC implementations.

	Without Battery	With Battery			
		Perfect Information	Stochastic $\rho = 0$	Deterministic $\rho = 0.1$	Value of Stochastic MPC
Expected Total Cost ($/month)	136,435	109,520	111,298	113,213	1,891
Expected Demand Charge ($/month)	-	-	136,397	136,433	-
Average FR Shortfall (kWh/month)	-	0	1,514	3,892	-
Value of Battery ($/month)	-	26,915	25,137	23,222	1,891
Improvement in Value of Battery	-	15.90%	8.24%	**Base**	-

3.3.3 *FR Shortfall and Constraint Violations*

We evaluated FR shortfall and constraint violations obtained with deterministic and stochastic MPC. We recall that the MPC controllers take no action when the policy becomes infeasible. This results in failing to deliver committed FR capacity, which could lead to penalties from the ISO. Figure 3.9 shows the probability and cumulative distributions for FR shortfall for the 200 validation scenarios and Figure 3.10 shows how often infeasibility occurs over the month. The results highlight that stochastic MPC without an FR buffer is more reliable at maintaining a feasible operation than deterministic MPC with a 10% buffer.

Motivated by the inability of deterministic MPC to capture FR signals using mean forecasts, we explored alternative deterministic representations for the forecast FR signal. Table 3.2 shows the results of using the minimum and maximum FR signals as forecasts. Both formulations lead to almost identical performance but their performance is worse than that achieved by the mean forecast (both in terms of cost and FR shortfall). We also created a *quasi-stochastic* MPC formulation that uses only three scenarios of the FR signal (including the mean value and minimum and maximum values) while mean forecasts were used for the rest of the disturbances. This approach can be interpreted as a specialized sampling technique that is targeted towards eliminating constraint violations. Notably, this

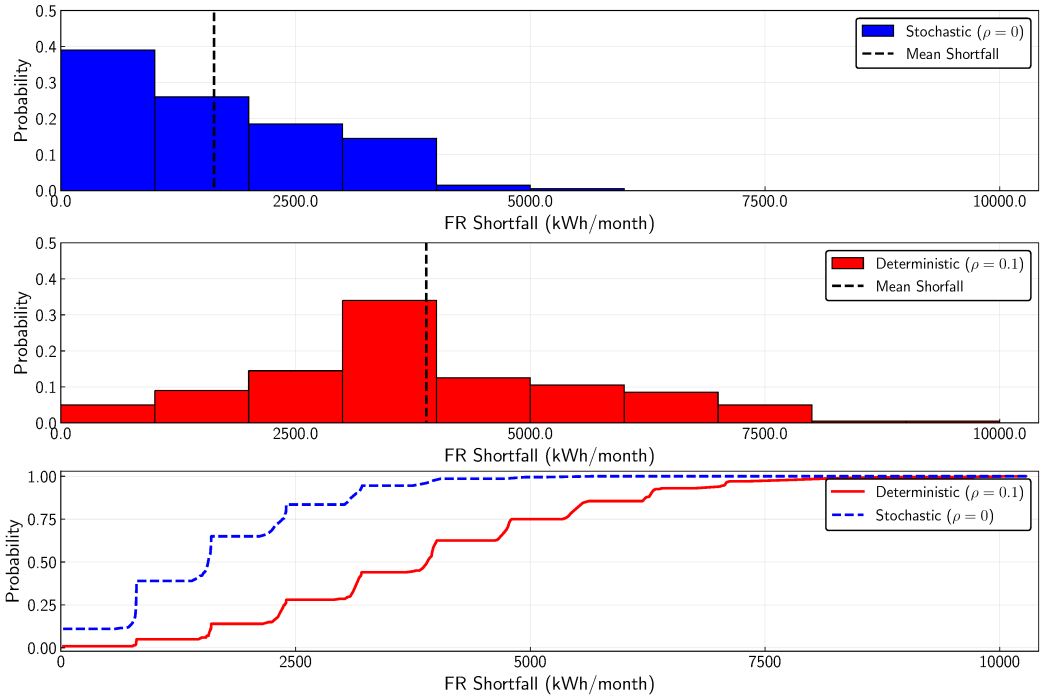

Figure 3.9: Probability and cumulative distributions for FR shortfall.

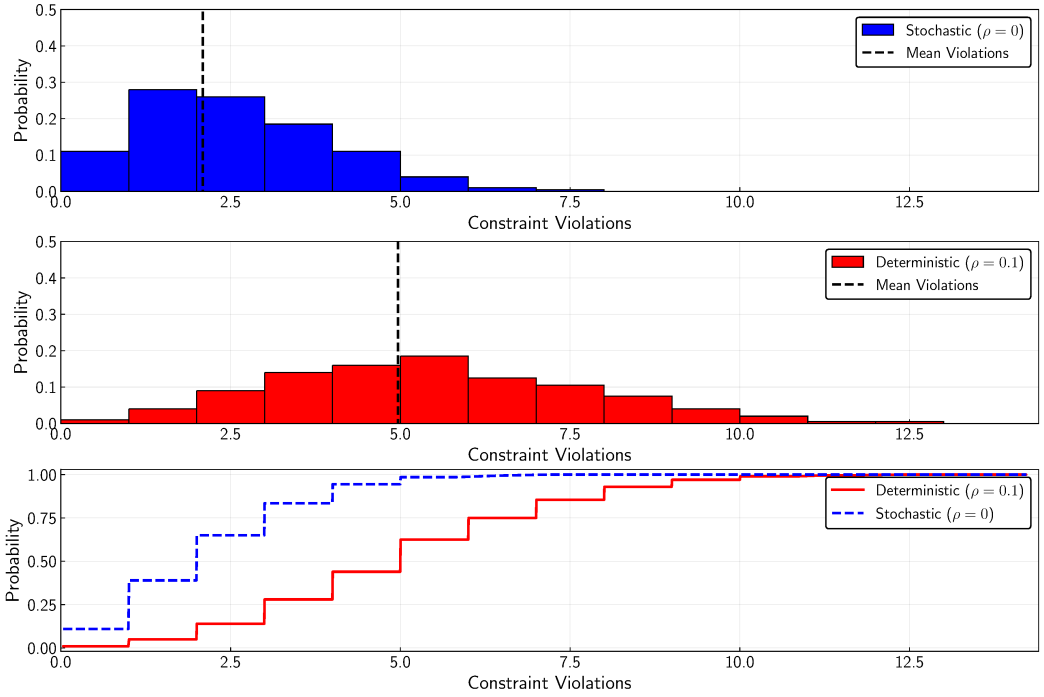

Figure 3.10: Probability and cumulative distributions for constraint violations.

simple formulation *fully eliminates constraint violations and FR shortfall* and achieves better economic performance than formulations with only minimum or maximum forecast FR signals. This highlights that constraint violations are indeed the result of using an inappropriate characterization of the FR signal. The quasi-stochastic approach, however, achieves worse economic performance than deterministic MPC (the value of the battery decreases by 13%). This indicates that the quasi-stochastic formulation is overly constrained (we use a buffer of $\rho = 0.1$), which restricts the full utilization of the battery capacity. In fact, in the results that we show next, we see that the use of a buffer in the quasi-stochastic MPC formulation is not necessary (because the three scenarios already capture the variability of the FR signal).

Table 3.2: Battery analysis using deterministic MPC with mean, minimum, and maximum forecasts for FR signal (with buffer of $\rho = 0.1$).

	Deterministic MPC $\mathbb{E}[\alpha]$	Deterministic MPC $\min(\alpha)$	Deterministic MPC $\max(\alpha)$	Quasi-Stochastic MPC $\mathbb{E}[\alpha], \min(\alpha), \max(\alpha)$
Expected Total Cost ($/month)	113,213	124,010	124,112	116,211
FR Revenue ($/month)	22,934	12,218	12,119	20,262
Average FR Shortfall (kWh/month)	3,892	84,384	88,505	0
Value of Battery ($/month)	23,222	12,425	12,323	20,224
Improvement in Value of Battery	**Base**	-46.50%	-46.93%	-12.91%

3.3.4 *FR Buffer Tuning*

We finally examine the effect of using different FR buffers on economic performance and on the amount of FR shortfall. Table 3.3 compares buffers of 0%, 5%, and 10% for both deterministic and stochastic MPC. The quasi-stochastic formulation is also studied with a buffer of 0%. As can be seen, deterministic MPC performs poorly without any buffer. The amount of FR shortfall is significant and this results in a higher total cost and lower battery value (because less FR revenue is realized). The amount of FR shortfall is also undesirable from the perspective of the ISO. Increasing the buffer leads to lower total cost

and higher battery value for deterministic MPC because the controller is able to participate more often by avoiding infeasibility issues. Increasing the buffer further eliminates the FR shortfall completely but increases cost and decreases the value of the battery (because the fraction of the battery capacity used for FR participation decreases). This inherent trade-off is shown in Table 3.4 and Figures 3.11 and 3.12. Here, we present a sensitivity analysis for deterministic MPC with different FR buffer values. Here, we can see that a buffer of 10% achieves the best cost and battery value. We note, however, that even a perfectly tuned deterministic MPC controller cannot achieve the performance of stochastic MPC (both in terms of economics and FR shortfall).

Table 3.3: Expected costs for MPC schemes with varying FR buffers.

Horizon	**7-day**						
Item	**Deterministic** ($\rho = 0$)	**Deterministic** ($\rho = 0.05$)	**Deterministic** ($\rho = 0.1$)	**Quasi-Stochastic** ($\rho = 0$)	**Stochastic** ($\rho = 0$)	**Stochastic** ($\rho = 0.05$)	**Stochastic** ($\rho = 0.1$)
Total Cost ($/month)	115,245	113,861	113,213	111,408	111,297	112,153	114,905
FR Revenue ($/month)	20,278	22,099	22,933	24,851	24,893	24,248	20,406
FR Shortfall (kWh/month)	103,049	24,658	3,892	1,582	1,514	152	0
Value of Battery ($/month)	21,190	22,574	23,222	25,027	25,137	24,282	21,530
Improvement in Value of Battery	-8.75%	-2.79%	**Base**	7.77%	8.25%	4.56%	-7.29%

Table 3.4: Expected costs and value of battery for deterministic MPC scheme with varying FR buffers.

Horizon	**7-day**					
Item	**Deterministic** ($\rho = 0$)	**Deterministic** ($\rho = 0.05$)	**Deterministic** ($\rho = 0.1$)	**Deterministic** ($\rho = 0.13$)	**Deterministic** ($\rho = 0.15$)	**Deterministic** ($\rho = 0.2$)
Total Cost ($/month)	115,245	113,861	113,213	113,398	113,778	116,757
FR Revenue ($/month)	20,278	22,099	22,933	22,851	22,545	19,709
FR Shortfall (kWh/month)	103,049	24,658	3,892	1,107	405	0
FR Shortfall Revenue ($/month)	4,581	1,092	172	51	17	0
Value of Battery ($/month)	21,190	22,574	23,222	23,037	22,657	19,678
Improvement in Value of Battery	-8.75%	-2.79%	**Base**	-0.80%	-2.43%	-15.26%

From Table 3.3 we also observe that the performance of the quasi-stochastic formula-

tion with 0% buffer is significantly better than that of the deterministic MPC formulations (for any buffer values). Notably, the FR shortfall for quasi-stochastic with no buffer is

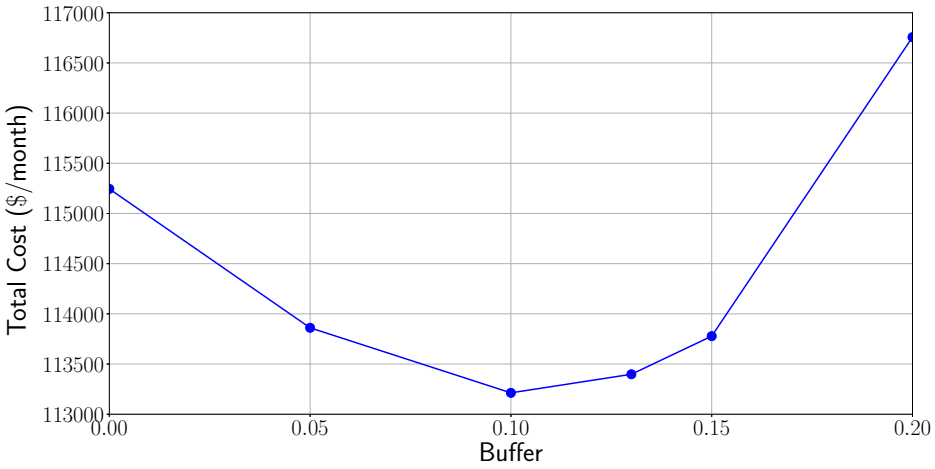

Figure 3.11: Economic performance (total cost) of deterministic MPC with varying FR buffers.

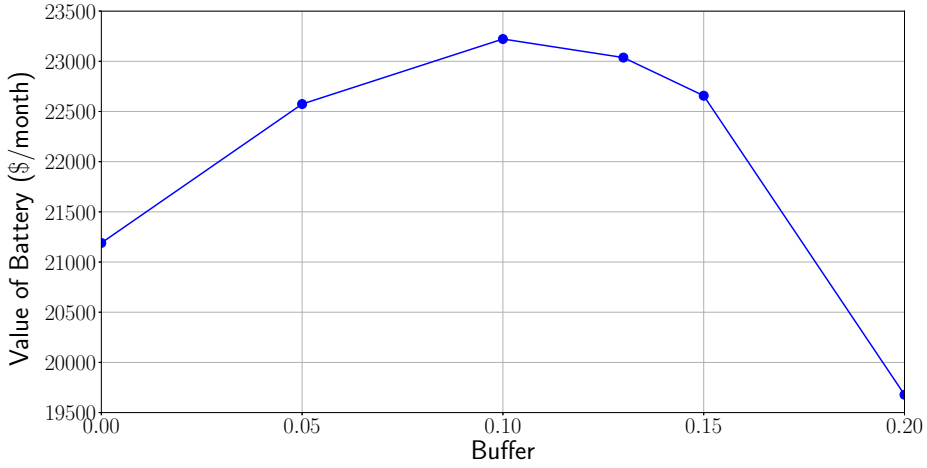

Figure 3.12: Economic performance (value of battery) of deterministic MPC with varying FR buffers.

similar to that achieved by stochastic MPC with no buffer (again reinforcing that the main source of infeasibility is the FR signal). We also see that the best economic performance is achieved with stochastic MPC using no buffer and that this achieves a smaller FR shortfall

than deterministic and quasi-stochastic MPC. This is because the stochastic MPC policy is successfully able to participate in the FR market and maintain battery feasibility. Increasing the buffer in stochastic MPC fully eliminates FR shortfall but this comes at the expense of a sharp decrease in economic performance (as the controller becomes more conservative).

3.4 Conclusions

In this chapter, we presented a computational framework that integrates forecasting, uncertainty quantification, and model predictive control (MPC) capabilities to provide a critical assessment of the performance of deterministic and stochastic MPC. By means of a battery management case study, we illustrate how off-the-shelf deterministic MPC implementations can suffer significant losses in performance and constraint violations due to their inability to handle disturbances that cannot be adequately represented by deterministic forecasts. Stochastic MPC provides a more systematic framework to handle these issues by directly capturing uncertainty descriptions of a wide range of disturbances.

4

STOCHASTIC MODEL PREDICTIVE CONTROL FOR CENTRAL HVAC PLANTS

In this chapter, we present a stochastic model predictive control (MPC) framework for central heating, ventilation, and air conditioning (HVAC) plants. The framework uses real data to forecast and quantify the uncertainty of disturbances affecting the system over multiple timescales (electrical loads, heating/cooling loads, and energy prices). We conduct detailed closed-loop simulations and systematic benchmarks for the central HVAC plant of a typical university campus. Results demonstrate that deterministic MPC fails to properly capture disturbances and that this translates into economic penalties associated with peak demand charges and constraint violations in thermal storage capacity (overflow and/or depletion) (Kumar et al., 2020). Our results also demonstrate that stochastic MPC provides a more systematic approach to mitigate uncertainties and that this ultimately leads to cost savings of up to 7.5% and the mitigation of storage constraint violations. Benchmark results also indicate that these savings are close to ideal savings (9.6%) obtained under MPC with perfect information.

4.1 Introduction

Commercial buildings are responsible for over 20% of the total energy consumption in the U.S. and annual expenditures of over $200 billion (Patel and Rawlings, 2019). In this context, heating, ventilation, and air-conditioning (HVAC) systems are the largest sources of energy use (nearly 50%) (DOE, 2011). Central HVAC plants are sophisticated systems that connect multiple energy carriers (water, electricity, natural gas, cooling water, hot water, steam) and equipment units (pumps, heat exchangers, cooling towers, chillers, and boilers) to meet the cooling and heating loads of single buildings or collections of buildings (e.g., university campuses and urban districts) (Rawlings et al., 2018). A central HVAC plant is the equivalent of a utility plant in a manufacturing facility. Besides total energy use, temporal profiles and peak use are also key factors that affect the efficiency and sustainability energy infrastructures. In particular, temporal profiles and peaks might push infrastructures to their design limits (e.g., capacity and ramping) and this might force operators to use inefficient back-up systems. Time-varying market prices and demand charges are used by operators and utility companies to try to mitigate such impacts. These pricing structures create an incentive for HVAC plants to incorporate thermal energy storage (TES) in order to shift loads in time and manipulate peak demands (Henze et al., 2004, 2008; Ma et al., 2012a; Risbeck et al., 2016, 2017). Effective operation of HVAC plants requires careful real-time management of the multiple components of the central plant; this is a challenging task because of the tight interconnection of equipment units, the presence of constraints, and the presence of multiple time-varying disturbances (e.g., energy loads and prices). Specifically, disturbances cannot be perfectly anticipated and thus complicate the planning process. All these factors are forcing commercial buildings to incorporate more sophisticated automation systems.

Model predictive control (MPC) is becoming an established automation technology in HVAC central plants (Avci et al., 2013; Deng et al., 2014; Oldewurtel et al., 2012; Mendoza-Serrano and Chmielewski, 2012; Kwadzogah et al., 2013; Rawlings et al., 2018). MPC can

anticipate and counteract disturbances and accommodate complex models, constraints, and cost functions (Rawlings et al., 2018; Kumar et al., 2019a; Qin and Badgwell, 2003). However, existing MPC implementations for HVAC central plants use deterministic representations of the disturbances. In particular, the most likely value (e.g., mean forecast obtained from autoregressive models) is used to compute control actions. The uncertainty associated with forecast errors is thus ignored during the computation of the control action and, instead, errors are counteracted through feedback. This deterministic approach is intuitive and works well in practice but might lead to cost degradation and failure to satisfy constraints (Magni and Scattolini, 2007; Limon et al., 2010). These issues are often overlooked in the MPC literature because benchmarking procedures often fail to systematically account for the effect of uncertainty (e.g., perfect forecasts are often assumed).

Uncertainty can be explicitly captured in the controller formulations such as stochastic MPC and robust MPC formulations (Kerrigan, 2001; Bernardini and Bemporad, 2009; Lucia et al., 2014; de la Penad et al., 2005; Mesbah, 2016). Robust MPC seeks to find optimal control actions to counteract extreme scenarios, whereas stochastic MPC seeks to determine the optimal actions by taking into consideration the probability of all possible occurrences. In a general multi-stage stochastic MPC, one assumes that uncertainty reveals progressively over time (at every stage) and this effect is modeled in the form of a scenario tree. This approach is intuitive as it captures how recourse would be implemented in an ideal setting but the scenario tree grows exponentially with the length of the prediction horizon (Lucia et al., 2014). The computational intractability of multi-stage MPC is often handled by using a two-stage approximation. Here, it is assumed that all uncertainty reveals after the first time stage and thus recourse is simplified. Some other alternatives for representing recourse include affine decision rules but these approaches tend to decrease flexibility (Farina et al., 2016; Paulson et al., 2017). In this work, we focus on a scenario-based two-stage stochastic MPC for HVAC central plants because we must consider long planning horizons.

In the context of energy systems, it has been recently reported that stochastic MPC can

systematically mitigate constraint violations and improve economic performance (Kumar et al., 2019a, 2018c). The benefits of stochastic MPC have also been widely reported in the context of building climate (airside) control (Oldewurtel et al., 2013; Zhang et al., 2013; Ma and Borrelli, 2012; Drgoňa et al., 2013) and energy management (Nghiem and Jones, 2017; Garifi et al., 2018; Ferrarini et al., 2014; Ma, 2012; Ma et al., 2014). We highlight that these studies have focused on the building (airside); to the best of our knowledge, stochastic MPC formulations for HVAC central plants have not been reported.

In this work, we present a computational framework for stochastic MPC for HVAC central plants. Our framework addresses HVAC plants for university campuses and seeks to assess the benefits of stochastic MPC over deterministic MPC. The framework uses real disturbance data to conduct forecasting and uncertainty quantification of disturbances. Our benchmarking procedure uses extensive closed-loop simulations under myriad realizations of disturbances in order to properly account for the effect of uncertainty in controller performance. Results indicate that deterministic MPC leads to violations of storage capacity constraints (overflow or drying up) of the hot and chilled water tanks and that stochastic MPC mitigates this issue. We find that storage capacity violations can be partially mitigated in deterministic MPC by adding buffer (back-off) terms but also that stochastic MPC consistently outperforms deterministic MPC in terms of cost. Specifically, we show that stochastic MPC achieves savings in the total cost of up to 7.52%. When these savings are disaggregated, we find consistent reductions in electricity cost (of 6.89%), in peak demand charges (of 29.8%), in natural gas cost (of 8.57%). We also find that stochastic MPC achieves significant reductions in natural gas usage and thus provides an effective approach to manipulate both electricity and natural gas demand profiles.

4.2 Computational Framework

In this section, we describe the computational framework used in our studies. We describe the decision-making setting, physical dynamic model, and disturbance forecast-

ing, and uncertainty quantification procedures. The framework incorporates deterministic, stochastic, and perfect information MPC formulations. The nomenclature used in this chapter is provided in Section A.2 of Appendix A.

4.2.1 *Decision-Making Setting*

The central HVAC plant for a typical university campus needs to produce chilled water and hot water in order to meet the time-varying loads (demands) from all buildings. The HVAC plant that we consider in this work consists of a chiller subplant that produces chilled water and a heat recovery (HR) chiller subplant that produces both chilled water and hot water, a hot water generator to produce hot water, cooling towers to reduce the temperature of the water purchased from the market, a dump heat exchanger (dump HX) for rejecting heat from the hot water, and storage tanks (one for chilled water and one for hot water). The goal is to determine hourly operating strategies for all equipment units so that the total cost of the external utilities that need to be purchased from the market (electricity, water, and natural gas) is minimized. Water and natural gas are charged on a total demand basis (at time-constant price) while total electricity is charged based on time-varying prices and the monthly peak electrical load is charged based on demand charges.

The various cost components faced by the central plant are:

- *Electricity transactions (hourly)*: The central plant purchases electricity required by the equipment for their operation. The transactions are charged at the time-varying market price, π_t^e.

- *Water transactions (hourly)*: The central plant needs to purchase water to make up for evaporative losses of water in the cooling towers. Water is purchased from the utility at a fixed price of π_t^w= \$0.009/gal.

- *Natural gas transactions (hourly)*: The central plant needs to purchase natural gas to run the hot water generator to satisfy the campus heating load. Natural gas is available from the utility at a fixed price of π_t^{ng}= \$0.018/kWh.

- *Peak electrical demand charges (monthly)*: The total electrical load (i.e., the central plant and attached campus load) is charged for the peak demand incurred over a month by the utility company (at a fixed demand charge price of π^D= \$4.5/kW).

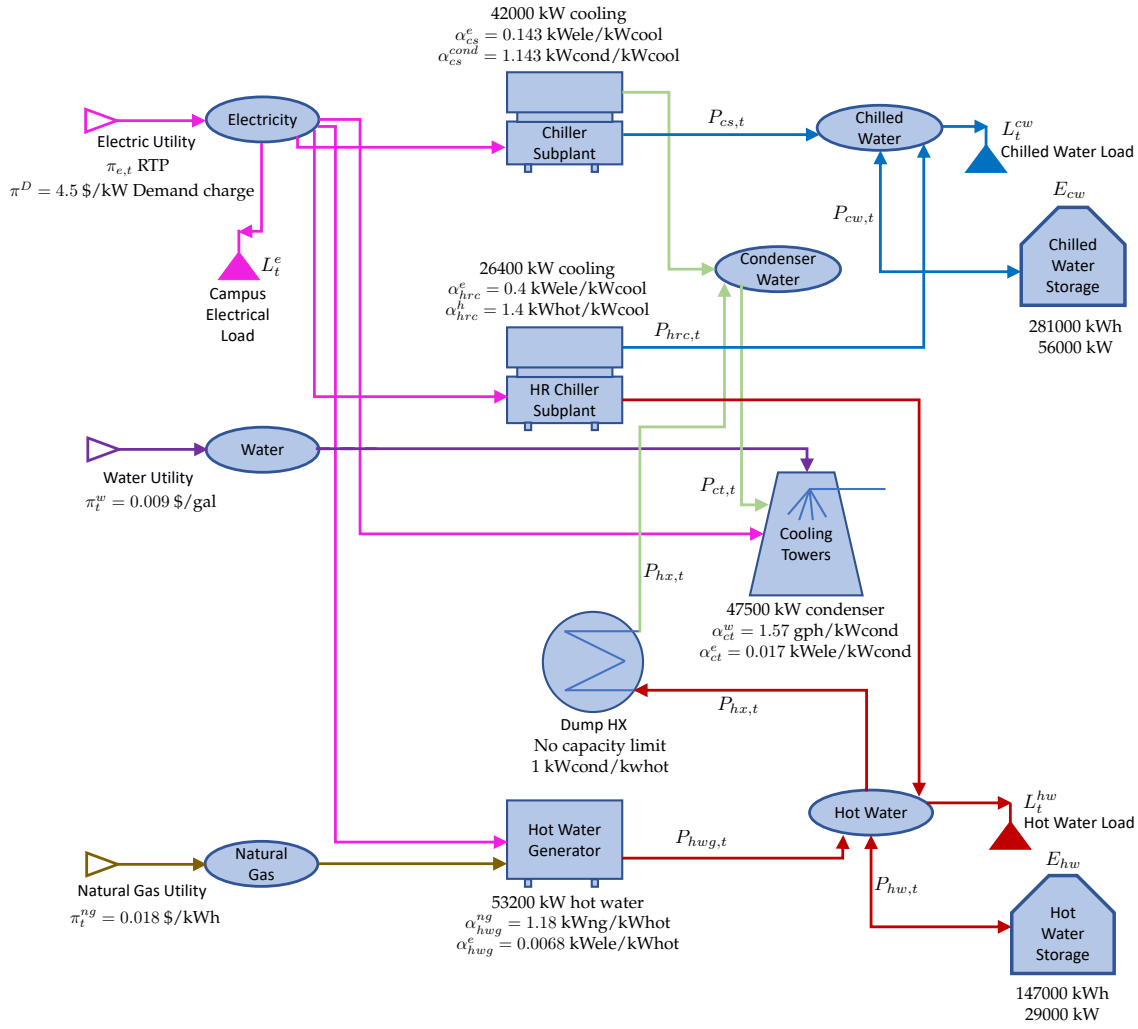

Figure 4.1: Schematic representation of central HVAC plant under study.

The HVAC central plant considered in this work is illustrated in Figure 4.1. This shows the energy flows between all the units and interactions with campus loads and utilities. The chiller subplant, HR chiller subplant, hot water generator, and cooling towers consume electricity in their operation. The cooling towers also consume utility water to make up for evaporative losses. The hot water generator is the only unit that consumes nat-

ural gas. The electricity, water, and natural gas consumption of these units are tied to their operating loads. The chiller subplant and HR chiller subplant consume α^e_{cs} and α^e_{hrc} kW of electricity per kW of chilled water produced, respectively; the hot water generator consumes α^e_{hwg} kW of electricity and α^{ng}_{hwg} kW of natural gas per kW hot water produced, respectively; and the cooling towers consume α^e_{ct} kW of electricity and α^w_{ct} utility water per kW of condenser water input, respectively.

The chilled water load (L^{cw}_t) of the campus is met by the chilled water production from chiller subplant ($P_{cs,t}$), the HR chiller subplant ($P_{hrc,t}$), and the discharge from chilled water storage ($P_{cw,t}$). The hot water load of the campus is met by hot water production from the HR chiller subplant ($\alpha^h_{hrc}P_{hrc,t}$), the hot water generator ($P_{hwg,t}$), and the discharge from the hot water storage ($P_{hw,t}$). The dump heat exchanger (HX) recycles excess hot water ($P_{hx,t}$) in the system by cooling it and producing condenser water which, together with the condenser water produced by the chiller subplant and HR chiller subplant, is cooled further by the cooling towers (total $P_{ct,t}$ condenser water is cooled by the towers). The manipulated variables for the system are the operating loads of all units, which include the chilled water production by the chiller and HR chiller subplants, hot water production by the hot water generator, discharge rates from the two storage tanks, the cooling load of the cooling towers and the heat exchange load of the dump HX.

The HVAC plant operations are driven by uncertain and time-varying disturbances, which are given by the campus loads for electricity (L^e_t), chilled water (L^{cw}_t), and hot water (L^{hw}_t), and by the electricity prices (π^e_t). The goal of the control (management) system of the plant is to determine operating loads for all units and storage levels to meet campus loads and to minimize the overall plant cost. Section A.2 of Appendix A provides a detailed description of all the variables and quantities involved.

4.2.2 *Forecasting and Uncertainty Quantification*

The proposed MPC formulations solve discrete-time optimal control problems at every time $t \in \mathbb{Z}_+$. These problems use a forecast (prediction) horizon of length N and denoted by the set $\mathcal{T} := \{t, t+1, t+2, \ldots, t+N-1\}$. Here, index t corresponds to the time instant $t \cdot h$, where h is the sampling time (assumed to be one hour). We implement zero-order hold for all quantities (these are held constant over the time interval $[(t-1) \cdot h, t \cdot h]$). We implement the MPC schemes under a receding-horizon framework (updated every sampling instant) over a simulation horizon of one year (denoted by the set $\mathcal{Y} := \{0, \ldots, Y\}$).

The control actions are the equipment operating loads while the states are the storage levels and other carryover quantities (e.g., peak demand). The decision to determine the operating loads is complicated by the fact that the time-varying electricity prices π_t^e and the campus electrical load, chilled water load, and hot water load exhibit uncertainty. The loads are denoted as $L_t = (L_t^e, L_t^{cw}, L_t^{hw})$. These disturbances are modeled as random variables with realizations indexed by ξ. At time t, a realization of the disturbances over a forecast horizon $\mathcal{T}$ is denoted as $L_{\mathcal{T}}(\xi), \pi_{\mathcal{T}}^e(\xi)$. The forecast trajectory used in deterministic MPC is a specific realization (usually the one with the highest probability) and is denoted as $\hat{L}_{\mathcal{T}}$ and $\hat{\pi}_{\mathcal{T}}^e$. To enable compact representation, we encapsulate all disturbances in the random vector $d_{\mathcal{T}} := (\pi_{\mathcal{T}}^e, L_{\mathcal{T}})$. We denote a realization of the disturbance $d_{\mathcal{T}}$ vector as $d_{\mathcal{T}}(\xi)$ and we denote the entire set of disturbance realizations using the notation $d_{\mathcal{T}}(\Xi)$.

We assume that the campus electrical load, chilled water load, hot water load, and electricity prices are Gaussian (normal) variables of the form $L_{\mathcal{T}} \sim \mathcal{N}(\hat{L}_{\mathcal{T}}, \Sigma_{\mathcal{T}}^L)$, and $\pi_{\mathcal{T}}^e \sim \mathcal{N}(\hat{\pi}_{\mathcal{T}}^e, \Sigma_{\mathcal{T}}^e)$. Time correlations in these disturbances are captured by the covariance matrices. Gaussian uncertainty descriptions are standard in autoregressive models and we have found them to provide satisfactory results in our studies. The mean and covariance matrices are updated using a receding-horizon scheme based on historical and real-time data (as it becomes available). We denote the observed (actual) disturbance history of H hours at time t as $d_{\mathcal{H}}$, where $\mathcal{H} := \{t-H, t-H+1, \ldots, t-1\}$.

In our implementation, the mean and covariance matrices for $L_{\mathcal{T}}$ and $\pi^e_{\mathcal{T}}$ are obtained using autoregressive (AR) models. Specifically, we use time series models of the form:

$$L_t = \sum_{k=1}^{q} \phi_k^L L_{t-k} + c^L + \epsilon_t^L \tag{4.2.1a}$$

$$\pi_t^e = \sum_{k=1}^{q} \phi_k^e \pi_{t-k}^e + c^e + \epsilon_t^e, \tag{4.2.1b}$$

where q is the order of the model; ϕ_t^L, ϕ_t^T, c^L and c^e are coefficients that are learned (estimated) from historical data; and ϵ_t^L and ϵ_t^e are white noise sequences. The mean forecasts $\hat{L}_{\mathcal{T}}$ and $\hat{\pi}^e_{\mathcal{T}}$ (the most likely realizations for Gaussian variables) are obtained by using the maximum likelihood estimates of the coefficients. Maximum likelihood procedures also provide estimates of the covariance matrices $\Sigma^L_{\mathcal{T}}$ and $\Sigma^e_{\mathcal{T}}$. Explicit techniques for performing maximum likelihood estimation are provided in Box et al. (2015). In this work, we use standard procedures provided in the `R` software package. We highlight that, in this work, we only consider measured disturbances and states. Consequently, no state estimation procedure is needed.

4.2.3 *Deterministic MPC*

The deterministic MPC controller uses the mean disturbance forecast $\hat{d}_{\mathcal{T}}$ to find the control policy. This policy minimizes the total cost that is forecast over the prediction horizon $\mathcal{T}$. This is done by solving an optimization problem at time $t \in \mathcal{Y}$. If the prediction horizon $\mathcal{T}$ at time t spans a month, the following optimization problem is solved (with the initial conditions provided by the already known current state values at time t, $E_{j,t}, ul_{j,t}, ol_{j,t}, j \in \{cw, hw\}$ and R_t):

$$\min \sum_{k\in\mathcal{T}} \sum_{j=\{e,w,ng\}} \hat{\pi}_k^j r_k^j + \frac{\pi^D}{\sigma_t} R_{t+1} + \sum_{k\in\mathcal{T}} \sum_{j\in\{cw,hw\}} \rho_j (ul_{j,k} + ol_{j,k}). \tag{4.2.2a}$$

$$\text{s.t. } r_k^e = \sum_{j\in\{cs,hrc,hwg,ct\}} \alpha_j^e P_{j,k} + \hat{L}_k^e,\ k \in \mathcal{T} \tag{4.2.2b}$$

$$r_k^j = \alpha_{i_j}^j P_{i_j,k},\ j \in \{w, ng\}, k \in \mathcal{T},\ i_w = ct, i_{ng} = hwg \tag{4.2.2c}$$

$$P_{ct,k} = \alpha_{cs}^{cond} P_{cs,k} + P_{hx,k},\ k \in \mathcal{T} \tag{4.2.2d}$$

$$P_{cs,k} + P_{hrc,k} + P_{cw,k} + S_{cw,k}^{un} - S_{cw,k}^{ov} = \hat{L}_k^{cw},\ k \in \mathcal{T} \tag{4.2.2e}$$

$$\alpha_{hrc}^h P_{hrc,k} + P_{hwg,k} - P_{hx,k} + P_{hw,k} + S_{hw,k}^{un} - S_{hw,k}^{ov} = \hat{L}_k^{hw},\ k \in \mathcal{T} \tag{4.2.2f}$$

$$E_{j,k+1} = E_{j,k} - P_{j,k},\ j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2g}$$

$$ul_{j,k+1} = ul_{j,k} - S_{j,k}^m,\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2h}$$

$$ol_{j,k+1} = ol_{j,k} - S_{j,k}^m,\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2i}$$

$$R_{t+1} \geq r_k^e \tag{4.2.2j}$$

$$R_{t+1} \geq R_t \tag{4.2.2k}$$

$$\underline{E}_{j,k} \leq E_{j,k} \leq \overline{E}_{j,k},\ j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2l}$$

$$\underline{P}_j \leq P_{j,k} \leq \overline{P}_j,\ j \in \{cs, hrc, hwg, ct, hx, cw, hw\}, k \in \mathcal{T} \tag{4.2.2m}$$

$$S_{j,k}^m \geq 0,\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2n}$$

$$ul_{j,k} \geq 0,\ j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2o}$$

$$ol_{j,k} \geq 0,\ j \in \{cw, hw\}, k \in \mathcal{T} \tag{4.2.2p}$$

The constraints (4.2.2b)-(4.2.2c) compute the demands of electricity, water, and natural gas (r_t^e, r_t^w, r_t^{ng}) that need to be purchased from the utility companies. Constraints (4.2.2d)-(4.2.2f) impose the energy balance for the condenser water. Constraints (4.2.2e) and (4.2.2f) are used to ensure that the chilled and hot water loads are met. Slack variables $S_{j,k}^m$, $m \in \{un, ov\}, j \in \{cw, hw\}$ are added in order to maintain feasibility in case of unmet (under-production) or overmet (over-production) of chilled water or hot water. The amounts of under-production or over-production of chilled and hot water are carried over using the state variables $ul_{j,k}$ and $ol_{j,k}$, $j \in \{cw, hw\}$. These amounts are computed using constraints (4.2.2h) and (4.2.2i). The unmet and overmet load state variables are penalized in the objective function by using the factors ρ_j, $j \in \{cw, hw\}$.

The constraints (4.2.2g) describe the dynamics of the state of charge (SOC) for chilled

and hot water TES. Constraint (4.2.2j) ensures that R_{t+1} is the peak residual electricity demand over the horizon $\mathcal{T}$. The parameter R_t is the carryover peak demand (the largest demand seen so far in the month). The peak residual demand is also considered a state of the system and is updated as $R_{t+1} = \max\{R_t, r_{t+1}(\xi)\}$. Constraints (4.2.2l)-(4.2.2p) provide bounds on the controls and states. The lower bounds for the operating loads of all units are zero except for the chilled water and hot water storage units (i.e., $\underline{P}_j = 0$ for $j \in \{cs, hrc, hwg, ct, hx\}$). For the storage units, the lower bounds of discharge rates correspond to the maximum charging rates, which are negative of the maximum discharging rates (i.e, $\underline{P}_j = -\overline{P}_j, j \in \{cw, hw\}$). The demand charge is weighted by a time-varying discounting factor $\sigma_t := \min\{(t_m - t)/N, 1\}$ (where t_m denotes the last hour of the month m). The discounting factor σ_t is used to adjust the demand charge because the prediction horizon N of the MPC formulation is shorter than a month (the period for which the peak electrical demand is charged). Also, it is desired to avoid the occurrence of a peak demand closer to the end of the month. Consequently, the time-varying factor $\sigma_t = \min\{(t_m - t)/N, 1\}$ is defined such that the demand charge is penalized higher when the current time reaches closer to the end of the month than that at the beginning of the month.

An important consideration in benchmarking of the performance of deterministic MPC is that the associated control policy might give rise to constraint violations in the storage levels (when the policy faces actual realizations of loads are observed in the time period $(t, t+1)$). This is because the control policy only satisfies constraints under the mean forecast. To avoid such constraint violations, the bounds on the chilled and hot water TES in constraints (4.2.2l) are modified to include a buffer capacity (a back-off term). A buffer value $\beta \in [0, 1]$ forces the storage capacity of chilled and hot water to vary between a fraction β and (1-β) of the maximum capacity (a value of $\beta = 0$ gives a standard MPC formulation). Such back-off terms are used to tune constraint violations under uncertain disturbances in applications of control and scheduling (Grosso et al., 2014; Koller et al., 2018; Kumar et al., 2018c). The buffer is updated after every sampling time in order to prevent the controller from being overly conservative. Specifically, when disturbances

push the system outside of the standard constraints, the bounds are adjusted until the optimizer restores feasibility (the following discussion provides more details).

Another important consideration when benchmarking performance is that the demand charge is accounted for based on the peak demand carried over at the end of the month. Consequently, the peak demand needs to be reset at the beginning of each month. To implement this over the simulation horizon $\mathcal{Y}$, we introduce the set $\mathcal{T}_M$ of ending time indices for all months. Index $t_m \in \mathcal{T}_M$ denotes the last hour of the month m. If the prediction horizon $\mathcal{T}$ at the current time t spans two months (i.e., $t < t_m$ and $t + N - 1 > t_m$ for some $t_m \in \mathcal{T}_M$), then the optimization problem at time t will consider peak demand variables $R_{1,t+1}$ and $R_{2,t+1}$ for the two months spanned. If at time t in the simulation, $t < t_m$ and $t + N - 1 > t_m$ for some $t_m \in \mathcal{T}_M$, we modify constraints (4.2.2j) and (4.2.2k) as follows:

$$R_{1,t+1} \geq r^e_k, \; k = t+1, t+2, \ldots, t_m \tag{4.2.3a}$$

$$R_{2,t+1} \geq r^e_k, \; k = t_m+1, t_m+2, \ldots, t+N-1 \tag{4.2.3b}$$

$$R_{i,t+1} \geq R_{i,t}, \; i = 1, 2. \tag{4.2.3c}$$

where $R_{2,t} = 0$ for $t = t_m$. The objective function (4.2.2a) is also modified to include the demand charges as:

$$\min \sum_{k\in\mathcal{T}} \sum_{j=\{e,w,ng\}} \pi^j_k r^j_k + \sum_{i=1}^{2} \frac{\pi^D_i}{\sigma_t} R_{i,t+1} + \sum_{k\in\mathcal{T}} \sum_{j\in\{cw,hw\}} \rho_j (ul_{j,k} + ol_{j,k}). \tag{4.2.4}$$

When $t > t_m + 1$ and $t < t_m + 2$, the optimization formulation (4.2.2a)-(4.2.2n) is solved (with $R_t = R_{2,t_{m+1}}$) because the prediction horizon spans only a single month ($m + 1$).

We denote the optimization problem solved in deterministic MPC at time t as $\mathcal{P}^{det}(x_t, \hat{d}_\mathcal{T})$. Here, the input arguments are the current state, control information, and disturbance forecast needed to solve the problem. We define the control action generated from the solution of the problem as $u_t = (P_{j,t}, j \in \{cs, hrc, hwg, ct, hx, cw, hw\})$ and the predicted states as $x_{t+1} = (E_{j,t+1}, ul_{j,t+1}, ol_{j,t+1}, R_{t+1}, j \in \{cw, hw\})$ (obtained assuming

forecast $\hat{d}_{t+1}$). The control action u_t is implemented in the system but the actual disturbance realized $d_{t+1}(\xi)$ will tend to deviate from the forecast $\hat{d}_{t+1}$. As a result, the state at time $t+1$ will differ from that predicted by the MPC controller. To account for the error in the prediction, the horizon is shifted forward to update the disturbance forecast using the new data history.

After obtaining the solution of $\mathcal{P}^{det}(x_t, \hat{d}_{\mathcal{T}})$, during the zero-order hold step (inter-optimization period), the current time decisions, u_t, for all units are held constant between $(t, t+1)$. During this inter-optimization period, the loads in real time can be varying, and the load balance constraints (4.2.2c) (production equal to consumption) will not hold. The loads will be balanced by charging or discharging the storage tanks as required in real time. Because the actual discharge does not exactly follow the discharge predicted by the optimization, the states will be at a different level than that predicted by the model at the end of zero-order hold. The amount of the difference is equal to the integrated difference between the constant discharge rate assumed by the optimization model and the actual time-varying discharge rate resulting from load following. We propose to model this integrated difference between the predicted discharge rate and the actual discharge rate by a normal random variable that is added to the predicted energy levels of the TES to simulate an actual energy level in the TES prior to solving the optimization problem on the next time step. The random variable is given by $v_{j,t} = \mathcal{N}(-0.5(L^j_{t+1} - L^j_t), 0.25\sigma^2_{j,err,t+1} + \sigma^2_{j,int})$, $j \in \{cw, hw\}$, where $\sigma^2_{j,err,t}$ is the variance of load prediction error for $t+1$ and $\sigma^2_{j,int}$ is the variance of integrated load for one-hour periods (obtained from historical data). The mean of the random variable $v_{j,t}$ is negative if the load rises in real time during the zero-order hold as the discharge from the TES will increase to make up for the higher load, and similarly, the mean of $v_{j,t}$ is positive if the load falls during the zero-order hold.

We summarize the receding-horizon scheme for deterministic MPC as:

1. START at $t = 0$ with the given $E_{j,0}$, and, $ul_{j,0} = 0, ol_{j,0} = 0$ for $j \in \{cw, hw\}$, $R_0 = 0$. Set $\underline{E}_{j,0} = \beta\overline{E}_j$, $\overline{E}_{j,0} = (1-\beta)\overline{E}_j$, $j \in \{cw, hw\}$. REPEAT for $t \in \mathcal{Y}$:

2. Use disturbance history $d_{\mathcal{H}}$ to obtain forecast $\hat{d}_{\mathcal{T}}$.

3. Solve $\mathcal{P}^{det}(x_t, \hat{d}_{\mathcal{T}})$ to obtain decisions $u_t = (P_{j,t}, j \in \{cs, hrc, hwg, ct, hx, cw, hw\})$.

4. Implement controls u_t over $(t, t+1)$.

5. Update storage states to actual states as $E_{j,t+1} = E_{j,t} - P_{j,t} + v_{j,t}, j \in \{cw, hw\}$.

6. Obtain the actual realized storage states by subtracting any unmet load from storage states obtained in the previous step: $E_{j,t+1} = E_{j,t+1} - ul_{j,t+1}, j \in \{cw, hw\}$.

7. Update the unmet and overmet load states to $ul_{j,t+1}, ol_{j,t+1}, j \in cw, hw$ from Eq. (4.2.2h) and (4.2.2i).

8. Modify the bounds on $E_{j,k}$ for $j \in \{cw, hw\}$ in constraints (4.2.2l) as follows:
   If $\beta\overline{E}_j \leq E_{j,t+1} \leq (1-\beta)\overline{E}_j$, set $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$.
   If $(1-\beta)\overline{E}_j \leq E_{j,t+1} \leq \overline{E}_j$, set $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = E_{j,t+1}$.
   If $0 \leq E_{j,t+1} \leq \beta\overline{E}_j$, set $\underline{E}_{j,t+1} = E_{j,t+1}$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$.
   If $E_{j,t+1} \geq \overline{E}_j$,set $E_{j,t+1} = \overline{E}_j$, $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = \overline{E}_j$, and update $ol_{j,k+1} = ol_{j,k+1} + (E_{j,t+1} - \overline{E}_j)$.
   If $E_{j,t+1} \leq 0$, set $E_{j,t+1} = 0$, $\underline{E}_{j,t+1} = 0$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$, and update $ul_{j,k+1} = ul_{j,k+1} - E_{j,t+1}$.

9. If the current prediction horizon $\mathcal{T}$ spans a single month, update the carryover demand charge $R_{t+1} = \max\{R_t, r_t^e\}$, else update as $R_{t+1} = (\max\{R_t, r_{t+1}^e\}, 0)$ with r_{t+1}^e being the actual realized residual electrical demand calculated from Eq. (4.2.2b) using actual *realized* electrical load $L_{t+1}^e(\xi)$.

10. Set $t \leftarrow t+1$. If $t = M$, STOP, otherwise RETURN to Step 2.

In this scheme, the actual realized disturbances are obtained from a set of validation scenarios $\tilde{\Xi}$. These scenarios are generated from actual disturbance data (not from the forecast). The performance of the MPC controller thus depends on the selection of the re-

alized disturbances and is thus random. We run the scheme for the entire set of validation samples to obtain probability distributions for diverse performance metrics such as cost.

4.2.4 *Stochastic MPC*

In stochastic MPC, uncertainty representations for disturbances are directly captured in the optimization formulation by using multiple realizations (scenarios) $d_{\mathcal{T}}(\xi)$, $\xi \in \bar{\Xi}$, where $\bar{\Xi}$ is a set of sample scenarios. The control decisions $u_t = (P_{j,t}, j \in \{cs, hrc, hwg, ct, hx, cw, hw\})$ for the next immediate hour (t+1) are here-and-now (commitment) decisions that need to be made prior to observing uncertainty. The controls $u_k(\xi) = (P_{j,k}(\xi), j \in \{cs, hrc, hwg, ct, hx, cw, hw\})$ at subsequent times $k \in \mathcal{T} \setminus \{t\}$ are modeled as wait-and-see (recourse) variables that can be corrected when the actual disturbance realization is observed. The stochastic MPC problem is thus formulated as a two-stage problem.

The residual and peak demands are also recourse decisions that are expressed as $r_{\mathcal{T}}(\xi)$ and $\max_{k\in\mathcal{T}} r_k(\xi)$. The SOC of the chilled water and hot water TES at time $t+1$ only depend on the previous storage $E_{j,t}$ and discharge rates $P_{j,t}$, $j \in \{cw, hw\}$. Consequently, $E_{j,t+1}, j \in \{cw, hw\}$ are also here-and-now variables. The rest of the trajectories $E_{j,k}(\xi), j \in \{cw, hw\}$ for $k \in \mathcal{T} \setminus \{t, t+1\}$ are recourse variables because the corresponding $P_{j,k}(\xi), j \in \{cw, hw\}$ are recourse variables for $k \in \mathcal{T} \setminus \{t\}$.

We use $\mathcal{P}^{sto}(x_t, d_{\mathcal{T}}(\bar{\Xi}))$ to denote the optimization problem solved in stochastic MPC at time t. The variables and constraints of the formulation are the same as those of the deterministic counterpart but are replicated for the set of realizations $\xi \in \bar{\Xi}$. We use non-anticipativity constraints to enforce the fact that the control actions u_t are here-and-now (commitment) decisions that need to be implemented in the system.

In the proposed stochastic MPC formulation, we solve a stochastic program at time $t \in \mathcal{Y}$ to minimize the expected total forecast cost of the system and satisfy the constraints under all scenarios over the prediction horizon $\mathcal{T}$. If the prediction horizon $\mathcal{T}$ at time t

spans a single month, the following problem is solved with the initial conditions provided by the already known current state values at time t, $E_{j,t}, ul_{j,t}, ol_{j,t}, j \in \{cw, hw\}$ and R_t:

$$\min \mathbb{E}\left[\sum_{k\in\mathcal{T}}\sum_{j=\{e,w,ng\}} \hat{\pi}^j_k(\xi) r^j_k(\xi) + \frac{\pi^D}{\sigma_t} R_{t+1}(\xi) + \sum_{k\in\mathcal{T}}\sum_{j\in\{cw,hw\}} \rho_j(ul_{j,k}(\xi) + ol_{j,k}(\xi))\right]. \tag{4.2.5a}$$

$$\text{s.t. } r^e_k(\xi) = \sum_{j\in\{cs,hrc,hwg,ct\}} \alpha^e_j P_{j,k}(\xi) + \hat{L}^e_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5b}$$

$$r^j_k(\xi) = \alpha^j_{i_j} P_{i_j,k}(\xi),\ j \in \{w, ng\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \text{where } i_w = ct, i_{ng} = hwg \tag{4.2.5c}$$

$$P_{ct,k}(\xi) = \alpha^{cond}_{cs} P_{cs,k}(\xi) + P_{hx,k}(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5d}$$

$$P_{cs,k}(\xi) + P_{hrc,k}(\xi) + P_{cw,k}(\xi) + S^{un}_{cw,k}(\xi) - S^{ov}_{cw,k}(\xi) = \hat{L}^{cw}_k(\xi),\ k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5e}$$

$$\alpha^h_{hrc} P_{hrc,k}(\xi) + P_{hwg,k}(\xi) - P_{hx,k}(\xi) + P_{hw,k}(\xi) + S^{un}_{hw,k}(\xi) - S^{ov}_{hw,k}(\xi) = \hat{L}^{hw}_k(\xi), \quad k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5f}$$

$$E_{j,k+1}(\xi) = E_{j,k} - P_{j,k}(\xi),\ j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5g}$$

$$ul_{j,k+1}(\xi) = ul_{j,k} - S^m_{j,k}(\xi),\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5h}$$

$$ol_{j,k+1}(\xi) = ol_{j,k} - S^m_{j,k}(\xi),\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5i}$$

$$R_{t+1}(\xi) \geq r^e_k(\xi),\ \xi \in \bar{\Xi} \tag{4.2.5j}$$

$$R_{t+1}(\xi) \geq R_t,\ \xi \in \bar{\Xi} \tag{4.2.5k}$$

$$\underline{E}_{j,k} \leq E_{j,k}(\xi) \leq \overline{E}_{j,k},\ j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5l}$$

$$\underline{P}_j \leq P_{j,k}(\xi) \leq \overline{P}_j,\ j \in \{cs, hrc, hwg, ct, hx, cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5m}$$

$$S^m_{j,k}(\xi) \geq 0,\ m \in \{un, ov\}, j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5n}$$

$$ul_{j,k}(\xi) \geq 0,\ j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5o}$$

$$ol_{j,k}(\xi) \geq 0,\ j \in \{cw, hw\}, k \in \mathcal{T}, \xi \in \bar{\Xi} \tag{4.2.5p}$$

$$P_{j,t}(\xi) = P_{j,t}(\xi'),\ j \in \{cs, hrc, hwg, ct, hx, cw, hw\}, \xi \neq \xi', \xi \in \bar{\Xi} \tag{4.2.5q}$$

The expected value $\mathbb{E}[\cdot]$ is defined over the set of scenarios $\bar{\Xi}$. We implement the

same approach as described for deterministic MPC for resetting the peak electrical demand charge at the beginning of each month, and for updating the bounds on the chilled water and hot water TES in every step of the MPC simulation. For consistency, here we present a stochastic MPC formulation that also accounts for a storage buffer β. In Section 4.3, we will see that stochastic MPC with no buffer ($\beta = 0$) can outperform deterministic MPC and significantly reduce constraint violations. This is because the controller naturally backs-off from the constraints when multiple disturbance realizations are accounted for.

We summarize the stochastic MPC scheme as:

1. START at $t = 0$ with the given $E_{j,0}$, and, $ul_{j,0} = 0, ol_{j,0} = 0$ for $j \in \{cw, hw\}$, $R_0 = 0$. Set $\underline{E}_{j,0} = \beta\overline{E}_j$, $\overline{E}_{j,0} = (1-\beta)\overline{E}_j$, $j \in \{cw, hw\}$. REPEAT for $t \in \mathcal{Y}$:

2. Use disturbance history $d_{\mathcal{H}}$ to obtain forecast $d_{\mathcal{T}}(\bar{\Xi})$.

3. Solve $\mathcal{P}^{sto}(x_t, d_{\mathcal{T}}(\bar{\Xi}))$ to obtain decisions $u_t = (P_{j,t}, j \in \{cs, hrc, hwg, ct, hx, cw, hw\})$.

4. Implement the decisions, u_t over $(t, t+1)$.

5. Update the storage states to the actual states as $E_{j,t+1} = E_{j,t} - P_{j,t} + v_{j,t}, j \in \{cw, hw\}$, where $v_{j,t}$ is the random variable as defined in Section 4.2.3.

6. Obtain the actual realized storage states by subtracting any unmet load from storage states obtained in the previous step: $E_{j,t+1} = E_{j,t+1} - ul_{j,t+1}, j \in \{cw, hw\}$.

7. Update the unmet and overmet load states to $ul_{j,t+1}, ol_{j,t+1}, j \in \{cw, hw\}$ from Eq. (4.2.5h) and (4.2.5i).

8. Modify the bounds on $E_{j,k}$ for $j \in \{cw, hw\}$ in constraints (4.2.5l) as follows:
   If $\beta\overline{E}_j \leq E_{j,t+1} \leq (1-\beta)\overline{E}_j$, set $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$.
   If $(1-\beta)\overline{E}_j \leq E_{j,t+1} \leq \overline{E}_j$, set $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = E_{j,t+1}$.
   If $0 \leq E_{j,t+1} \leq \beta\overline{E}_j$, set $\underline{E}_{j,t+1} = E_{j,t+1}$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$.
   If $E_{j,t+1} \geq \overline{E}_j$, set $E_{j,t+1} = \overline{E}_j$, $\underline{E}_{j,t+1} = \beta\overline{E}_j$, $\overline{E}_{j,t+1} = \overline{E}_j$, and update $ol_{j,k+1} = ol_{j,k+1} + (E_{j,t+1} - \overline{E}_j)$.

If $E_{j,t+1} \leq 0$, set $E_{j,t+1} = 0$, $\underline{E}_{j,t+1} = 0$, $\overline{E}_{j,t+1} = (1-\beta)\overline{E}_j$, and update $ul_{j,k+1} = ul_{j,k+1} - E_{j,t+1}$.

9. If the current prediction horizon $\mathcal{T}$ spans a single month, update the carryover demand charge $R_{t+1} = \max\{R_t, r_t^e\}$, else update as $R_{t+1} = (\max\{R_t, r_{t+1}^e\}, 0)$ with r_{t+1}^e being the actual realized residual electrical demand calculated from Eq. (4.2.5b) using actual *realized* electrical load $L_{t+1}^e(\xi)$.

10. Set $t \leftarrow t+1$. If $t = M$ STOP, otherwise RETURN to Step 2.

In this scheme, the actual realized disturbances are obtained from the set of validation samples $\tilde{\Xi}$. Importantly, these validation samples differ from the realizations used in the MPC controller formulation $\bar{\Xi}$. By running the stochastic MPC scheme for all validation samples, we can compute probability distributions for performance metrics that are compared with those from deterministic MPC. This systematic procedure ensures fair comparisons between different MPC implementations.

4.2.5 *Perfect Information MPC*

We also consider a perfect information MPC implementation to evaluate the ideal performance of MPC. Under perfect information MPC, we compute commitment policies $u_t(\xi)$ at every time t for every realization $\xi \in \tilde{\Xi}$ of the loads and prices. These policies can be computed by removing the non-anticipativity constraints (4.2.5q) from the stochastic MPC formulation (4.2.5) and replacing the disturbance forecast model with the true disturbance signals corresponding to each realization $d(\xi)$, and implementing the same scheme as the stochastic MPC run over a year period $\mathcal{Y}$ (Kumar et al., 2018c).

4.2.6 *Handling Constraint Violations*

The deterministic MPC controller can violate the constraints during the transition $(t, t+1)$ if the forecasts are poor or if the scenarios used. The stochastic MPC formulation can also

incur violations if the scenarios used do not capture the actual realizations. To capture this issue in our closed-loop simulations, an auxiliary MPC controller is used to correct the control actions and restore feasibility. At time t, this controller solves the feasibility restoration problem:

$$\min_{\substack{\Delta P_j, \\ j\in\{cs,hrc,hwg,ct,hx,cw,hw\}}} \sum_j |\Delta P_j| \tag{4.2.6a}$$

$$\text{s.t. } E_{j,t+1} = E_{j,t} - P_{j,t} + \Delta P_j,\ j \in \{cw, hw\} \tag{4.2.6b}$$

$$\sum_{j\in\{cs,hrc,cw\}} (P_{j,t} + \Delta P_{j,t}) = L_t^{cw} \tag{4.2.6c}$$

$$\sum_{j\in\{hwg,hw\}} (P_{j,t} + \Delta P_j) + \alpha_{hrc}^h \left(P_{hrc,t} + \Delta P_{hrc}\right) - P_{hx,t} - \Delta P_{hx} = L_t^{hw} \tag{4.2.6d}$$

$$0 \le E_{j,t+1} \le \overline{E}_j,\ j \in \{cw, hw\} \tag{4.2.6e}$$

$$\underline{P}_j \le P_{j,t} + \Delta P_j \le \overline{P}_j,\ j \in \{cs, hrc, hwg, ct, hx, cw, hw\}. \tag{4.2.6f}$$

This formulation uses the actual realizations for the chilled water and hot water loads over the time interval $(t, t+1)$. The feasibility restoration problem seeks to find a net rate correction $\Delta P_j, j \in \{cs, hrc, hwg, ct, hx, cw, hw\}$ that satisfies the storage constraints. If the auxiliary controller fails to remain feasible even after solving the restoration problem, we assume that no action is taken. In other words, we set $P_{j,t} = 0$ for $j \in \{cs, hrc, hwg, ct, hx, cw, hw\}$ and correct the states $E_{j,t+1}$ for $j \in \{cw, hw\}$ and R_{t+1} accordingly. This leads to loss of performance because there is a failure to meet the campus loads during time $(t, t+1)$.

4.2.7 *Benchmarking Procedure*

We extend the benchmarking procedure given in Kumar et al. (2019a) for battery management systems. To distinguish the policies obtained from the three MPC schemes, we denote the policies obtained from the deterministic MPC as u_t^{det} and $r_t^{det}(\xi) = (r_t^{j,det}, j \in \{e, w, ng\})$, those from stochastic MPC scheme as u_t^{sto} and $r_t^{sto}(\xi) = (r_t^{j,sto}, j \in \{e, w, ng\})$,

and those from perfect information MPC as $u_t^{perf}(\xi)$ and $r_t^{perf}(\xi) = (r_t^{j,perf}, j \in \{e, w, ng\})$ (corresponding to each realization ξ).

Each realization in the validation set $\tilde{\Xi}$ generates an annual cost for the MPC controllers. The annual cost under stochastic MPC for a given realization $\xi \in \tilde{\Xi}$ and under a given closed-loop policy $u_{\mathcal{M}}, r_{\mathcal{M}}$ is given by:

$$\Phi^{sto}(\xi) := \sum_{t \in \mathcal{Y}} \sum_{j=\{e,w,ng\}} \pi_k^j(\xi) r_k^{j,sto}(\xi) + \sum_{m=1}^{12} \pi^D \max_{t \in \{1,\ldots,t_m\}} r_t^{e,sto}(\xi). \tag{4.2.7}$$

The annual cost for deterministic MPC is denoted as $\Phi^{det}(\xi)$ and is defined as in (4.2.7) (but with its corresponding closed-loop policy u_t^{det} and $r_t^{det}(\xi) = (r_t^{j,det}, j \in \{e, w, ng\})$). The annual cost for perfect information MPC is:

$$\Phi^{perf}(\xi) := \sum_{t \in \mathcal{Y}} \sum_{j=\{e,w,ng\}} \pi_k^j(\xi) r_k^{j,perf}(\xi) + \sum_{m=1}^{12} \pi^D \max_{t \in \{1,\ldots,t_m\}} r_t^{e,perf}(\xi). \tag{4.2.8}$$

The costs for the different validation realizations are used to create empirical probability distributions and cumulative probability distributions for diverse quantities of interest and to compute statistics such as expected costs $\mathbb{E}[\Phi^{det}(\tilde{\Xi})]$, $\mathbb{E}[\Phi^{sto}(\tilde{\Xi})]$, and $\mathbb{E}[\Phi^{perf}(\tilde{\Xi})]$.

Of particular interest in our benchmark studies is a metric that we call the *expected cost of the HVAC central plant*. To compute this value, we evaluate the total cost under the assumption that there is no HVAC central plant serving the campus. For a particular validation realization $\xi \in \tilde{\Xi}$, this cost is denoted as $\Phi^{nocp}(\xi)$. The ideal expected cost of the central plant is defined as $\text{CCP}^{perf}(\xi) := \Phi^{perf}(\xi) - \Phi^{nocp}(\xi)$. The cost of the central plant under stochastic MPC is $\text{CCP}^{sto}(\xi) := \Phi^{sto}(\xi) - \Phi^{nocp}(\xi)$ and under deterministic MPC is $\text{CCP}^{det}(\xi) := \Phi^{det}(\xi) - \Phi^{nocp}(\xi)$. As in the case of cost, the realizations are used to obtain probability distributions and to compute statistics such as $\mathbb{E}[\text{CCP}^{perf}(\tilde{\Xi})]$, $\mathbb{E}[\text{CCP}^{sto}(\tilde{\Xi})]$, and $\mathbb{E}[\text{CCP}^{det}(\tilde{\Xi})]$. The cost of the central plant is a metric that reflects losses/gains in asset value due to the use of better control policies (it isolates the effect of the control from that of the equipment). We also consider the value of the stochastic MPC, which is defined

as $\text{VSMPC}(\xi) := \text{CCP}^{det}(\xi) - \text{CCP}^{sto}(\xi)$, and the expected value of stochastic MPC as $\mathbb{E}[\text{VSMPC}(\tilde{\Xi})]$.

4.3 Benchmark Results

We now present closed-loop simulation results for deterministic, stochastic, and perfect information MPC for an entire year of operation. The controllers use a prediction horizon of 168 hours (seven days), which is chosen based on the observation that the data for the loads and electricity prices exhibit weekly periodicity (e.g., high load on weekdays and low load on weekends). In other words, a horizon of seven days captures periodic effects (Kumar et al., 2018c). The stochastic MPC problem contains 100 forecast scenarios and a total of 168,450 variables and 143,750 constraints (the realizations are obtained using Monte Carlo sampling). A total of 200 closed-loop year-long runs were performed for each MPC implementation (using the same validation scenarios). The number of forecast and validation scenarios are chosen considering a trade-off between performance and computational cost. Specifically, using a few scenarios leads to fast computations but the full uncertainty space is not well represented and this limits the benefits of stochastic MPC. Having a larger number of scenarios increases the computational time but the controller performance improves. Hence, we use a heuristic-based approach to chose a suitable number of forecast and validation scenarios to benchmark stochastic MPC. The simulations were run on a 32-core machine with Ubuntu 14.04, Intel Xeon 2.30 GHz processors, and 188 GB RAM. The schemes are implemented in `Julia` and leverage the algebraic modeling capabilities of `JuMP` (Dunning et al., 2017). The optimization problems are solved in extensive form using `Gurobi 8.1`.

We use time series forecasting procedures provided in the `R` software package. Specifically, we use the `ar` function in `R` to estimate the coefficients of the AR model of order q with the settings as `aic=false`, `order=`q, and `method="ols"` (ordinary least-squares). The forecasts are obtained using the `forecast` function with a prediction horizon of N,

and covariance matrices for the forecasts are obtained using `var.pred` function. With this information, we generate a set of disturbance realizations $\bar{\Xi}$ by sampling from the corresponding estimated probability density functions. Forecasts and scenarios for the electrical load, chilled water load, hot water load, and the electricity price are obtained using a q =168 order AR model that predicts the loads over a horizon of N =168 hours. The parameters of the AR models are estimated by using a data history of 184 days (i.e., 6 months) at every time step in the closed-loop of MPC. We highlight that an order $q = 168$ is important to preserve the strong weekly periodicity of the data in the AR prediction model. If an AR model of order less than 168 is used, the 1-week long forecast from this model is unable to properly capture the weekly periodicity (e.g. weekdays and weekends) of the data and results in high prediction errors (e.g., such a model has high prediction errors during weekends if order $q = 96$, equivalent to 4 days, is used).

Each MPC problem instance takes, on average, about one second to solve for deterministic MPC and about 6-7 seconds to solve for stochastic MPC. Despite these fast solutions, we note that year-long closed-loop simulations required approximately *2 hours* for deterministic MPC and *10 hours* for stochastic MPC. These computational times comprise forecasting, optimization solution, and feasibility checks. These computational times are serial but can be partially parallelized (this is left as a topic of future work since the computational workflows involved are complex).

4.3.1 *Forecasting of Loads and Electricity Prices*

Figures 4.2-4.4 show historical data for the campus electrical load, hot water load, and chilled water load for the entire year. Figure 4.5 shows historical electricity prices for the same period. The vertical red lines represent monthly periods.

A single instance for the 1-week forecasts for electrical load, cold water load, hot water load, and electricity price is shown in Figures 4.6, 4.7, 4.8, and 4.9, respectively. The mean forecasts are represented by the dark bold curves and the 99% confidence intervals are

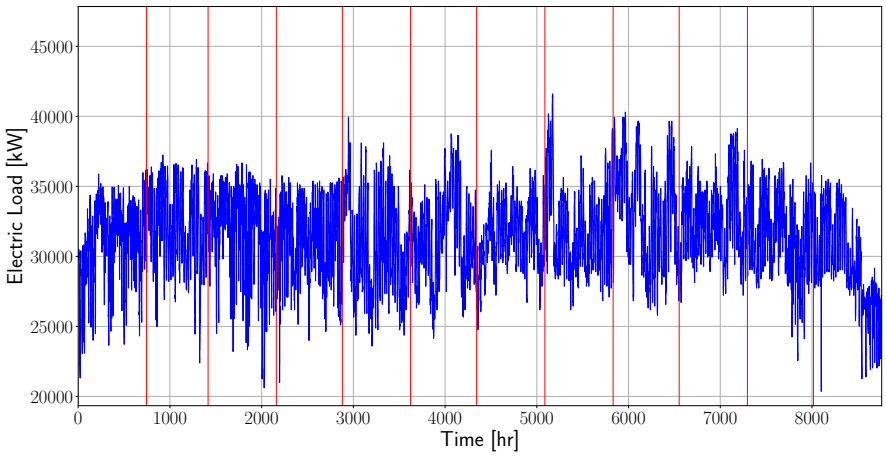

Figure 4.2: Historical electrical load of the campus. Red vertical lines denote end of each month.

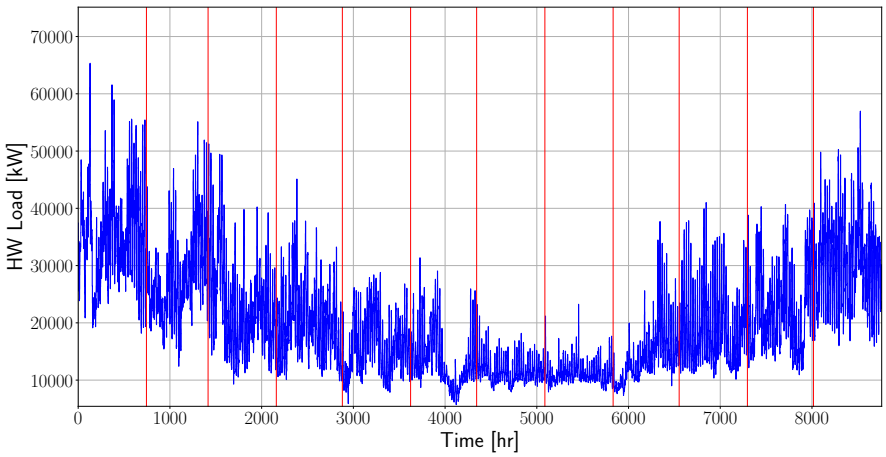

Figure 4.3: Historical hot water load of the campus. Red vertical lines denote end of each month.

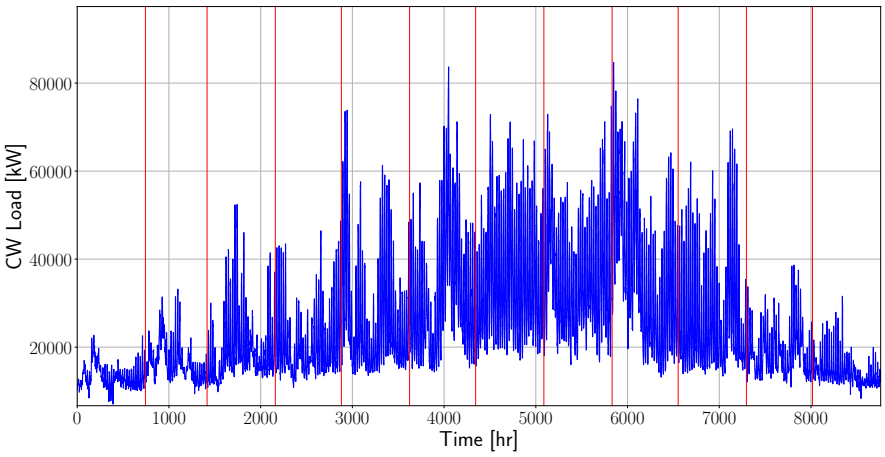

Figure 4.4: Historical chilled water load of the campus. Red vertical lines denote end of each month.

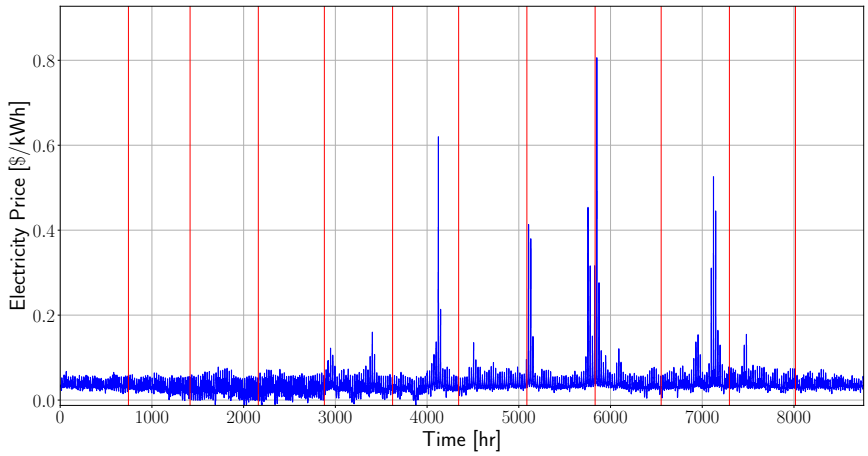

Figure 4.5: Historical electricity price data. Red vertical lines denote end of each month.

shown in grey color. Sample scenarios are shown as light black curves. From these results, we see that the AR models can capture the trends of the disturbances but that significant uncertainty exists. In particular, we notice that the magnitude of the confidence interval (the range) rises sharply within the first few hours. As we show next, this will be a major factor that drives constraint violations.

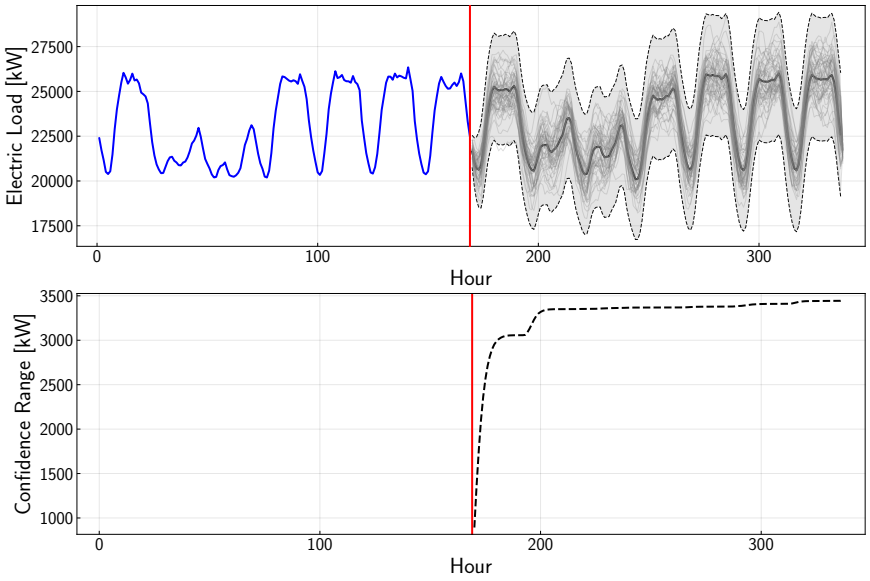

Figure 4.6: A single instance of electrical load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time.

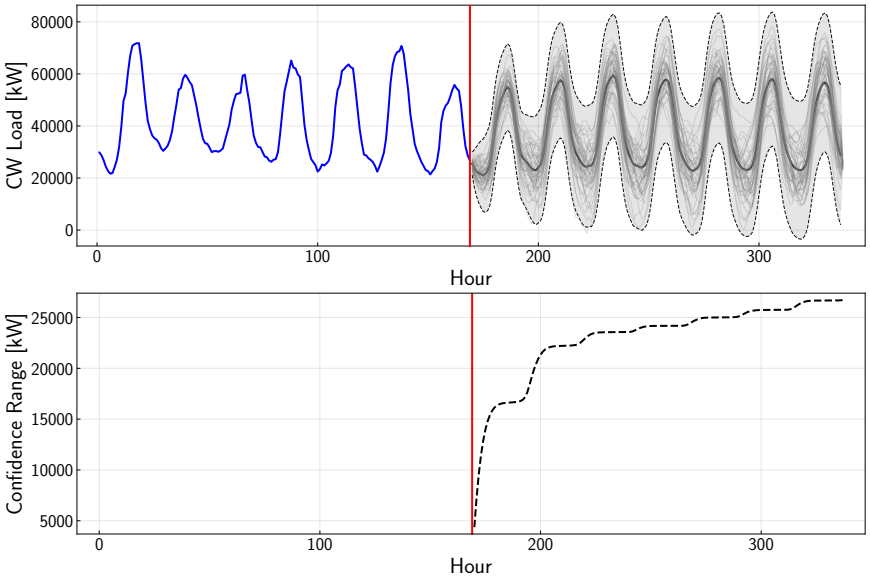

Figure 4.7: A single instance of chilled water load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time.

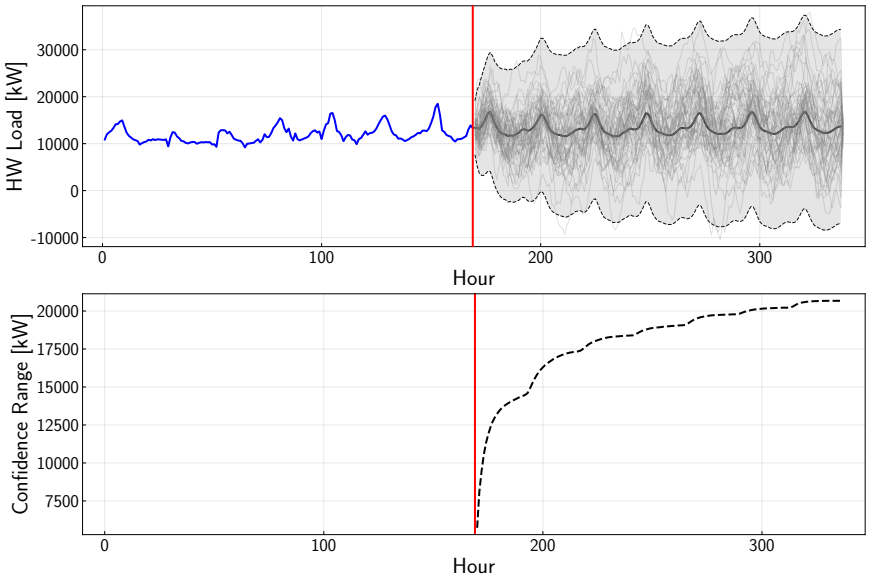

Figure 4.8: A single instance of hot water load forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time.

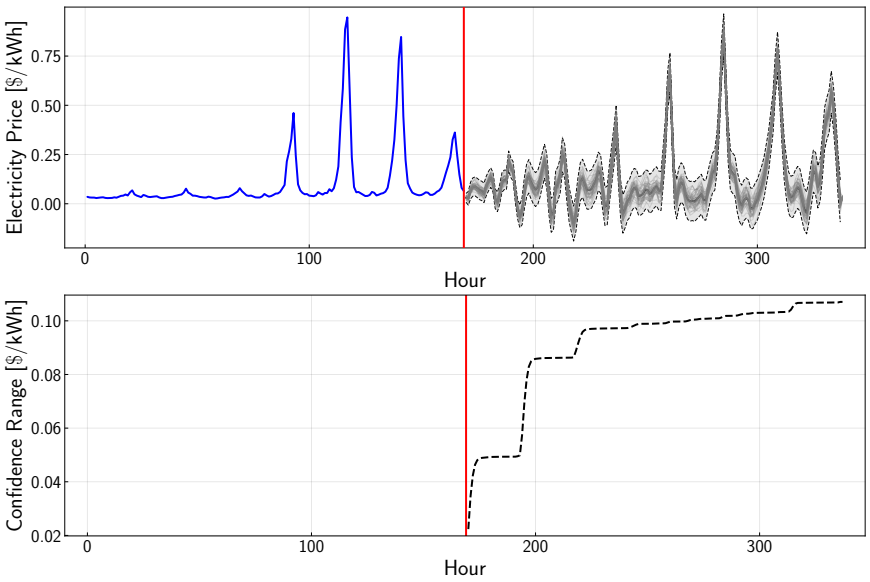

Figure 4.9: A single instance of electricity price forecast for 1-week with AR model. Vertical red line denotes the current time. In the top panel, dark bold curve represents mean forecast, the grey band denotes 99% confidence interval and the light black curves within the band represent a few sample scenarios. The bottom panel shows the trajectory of the 99% confidence range with prediction time.

4.3.2 *Closed-Loop Performance*

We compare the closed-loop policies for deterministic and stochastic MPC. In these results, a storage buffer of $\beta = 0$ is used for stochastic MPC and a buffer of 10% ($\beta = 0.1$) is used for deterministic MPC. Figure 4.10 provides a snapshot for a given validation scenario for deterministic MPC. Here, we note that the controller uses the storage buffer fairly frequently to counteract uncertainty in the disturbances. Figure 4.11 shows a snapshot for stochastic MPC under a given validation scenario. Here, we note that the controller does not require an explicit buffer for storage, which results in better utilization of the storage. An animation of the closed-loop performance of deterministic and stochastic MPC can be found at `https://github.com/zavalab/JuliaBox/tree/master/HVAC_Plant` (these help visualize the closed-loop dynamics).

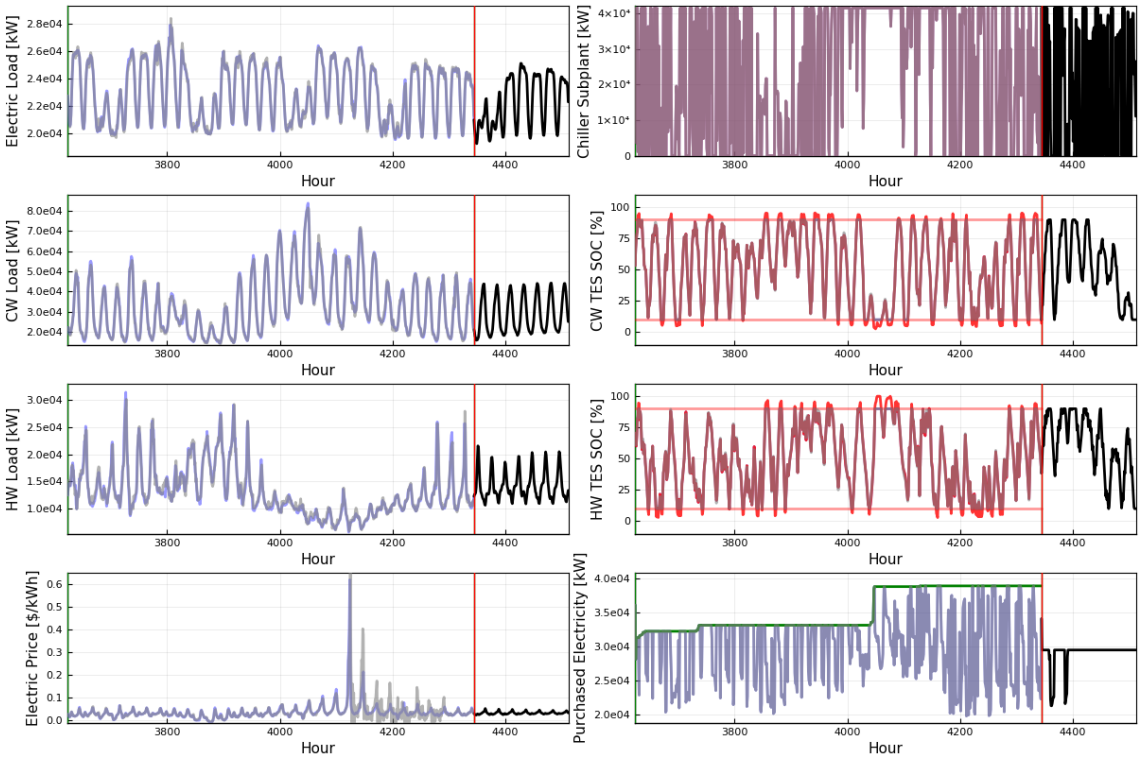

Figure 4.10: Closed-loop profile for deterministic MPC with $\beta = 0.1$. Black lines represent forecasts and model predictions. Blue lines represent actual realizations. For the control policies, red lines represent the actual implemented policy. For the state of charge, red horizontal lines represent the storage buffer. For the residual electrical load, the green line represents the running peak R_t.

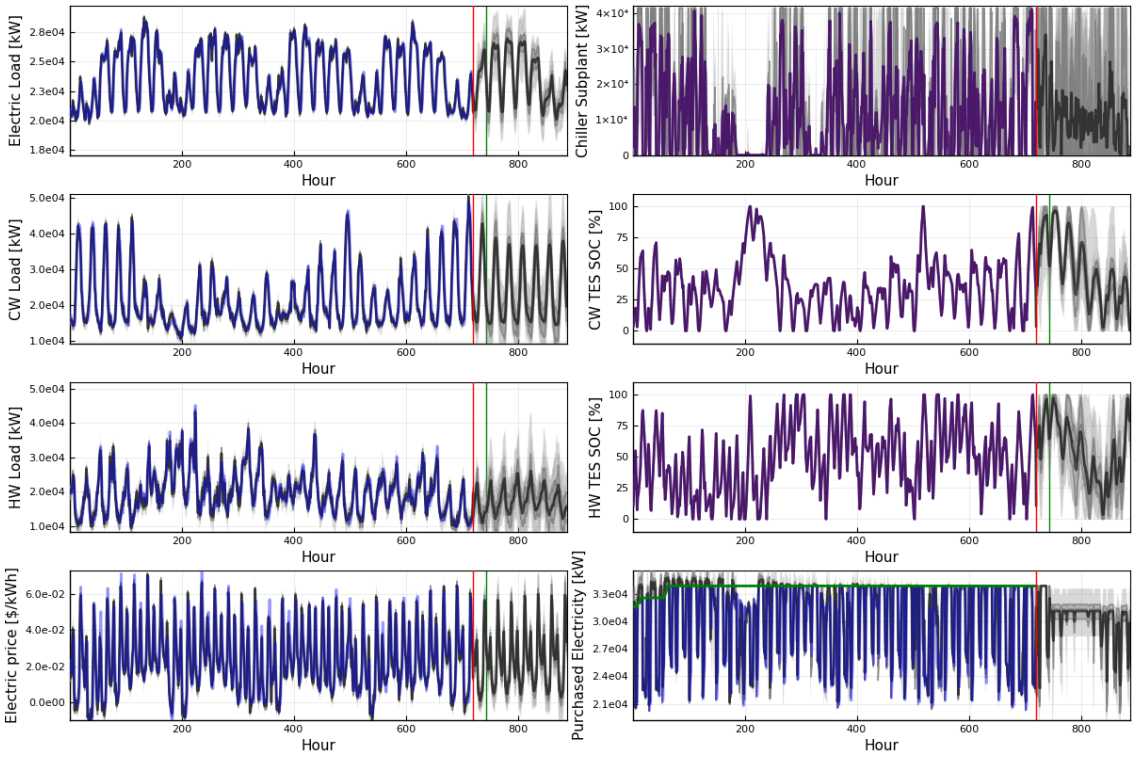

Figure 4.11: Closed-loop profile for stochastic MPC with $\beta = 0$. The grey regions represent uncertainty forecasts with black representing mean forecasts. Blue lines represent realized observations or committed policies. For the control policies, red lines represent the actual implemented policy and usually overlap with committed policy. For the residual load, the green line represents the peak observed residual electricity demand.

4.3.3 *Economic Performance and Constraint Violations*

In Table 4.1 we compare utility usage of the campus together with the HVAC central plant under the MPC implementations. We observe that the system would ideally (under perfect information) consume 264,800 MWh of electricity, 197,839 million gallons of water, and 29,067 MWh of natural gas in the year. These numbers represent the best possible performance under a finite horizon MPC and highlight the high resource use of these systems. The system under a deterministic MPC implementation has a higher consumption for each of the utilities. This is expected because deterministic MPC faces forecast errors. The stochastic MPC implementation reduces natural gas use of the deterministic counter-

Table 4.1: Utility usage analysis for different MPC implementations.

Expected Utility Usage	Campus With Central Plant			
	Perfect Information	Stochastic $\beta = 0$	Deterministic $\beta = 0.1$	Savings by Stochastic MPC
Electricity (MWh/year)	264,800 **(-0.2%)**	265,137 **(-0.07%)**	265,324 **(Base)**	187
Water (MMgal/year)	197,839 **(-1.8%)**	199,593 **(-0.96%)**	201,520 **(Base)**	1,927
Natural Gas (MWh/year)	29,067 **(-12.3%)**	30,315 **(-8.57%)**	33,155 **(Base)**	2,840

part by a significant amount. Specifically, under stochastic MPC, we observe reductions in natural gas use of 8.57% (2,840 MWh). These results highlight that stochastic MPC can reduce the resource footprint by better handling of disturbances. Stochastic MPC achieves modest reductions in electricity and water consumption of 0.07% (187 MWh) and 0.96% (1,927 million gallons), respectively. We will see, however, that significant cost reductions are achieved (due to the time-varying nature of electricity prices).

Table 4.2: Economic performance analysis for different MPC implementations (campus with central plant).

	Campus Only (No Central Plant)	**Campus With Central Plant**			
		Perfect Information	Stochastic $\beta = 0$	Deterministic $\beta = 0.1$	Value of Stochastic MPC
Expected Electricity Cost ($/year)	9,764,251	11,370,066	11,422,586	11,545,022	122,436
Expected Water Cost ($/year)	-	1,780,554	1,796,342	1,813,681	17,339
Expected Natural Gas Cost ($/year)	-	523,210	545,679	596,801	51,122
Expected Demand Charge ($/year)	2,051,316	2,273,331	2,342,457	2,500,522	133,569
Expected Total Cost ($/year)	11,815,567	16,047,162	16,107,064	16,456,027	348,963
Cost of Central Plant ($/year)	-	4,231,595	4,291,497	4,640,460	348,963
Savings in Central Plant Cost	-	9.66%	7.52%	**Base**	7.52%

The economic performance of the MPC implementations is summarized in Table 4.2. Here, we present the expected total cost and we disaggregate this cost in its different components. The expected total cost of the campus (without the central plant) is 11,815,567 $/year and the expected total cost for the campus with the central plant (operated with perfect information MPC) is 16,047,162 $/year. This indicates that the operation of the HVAC central plant alone costs 4,231,595 $/year (this shows the large costs associated with the plant). We factor out the cost of the central plant from the total cost because the MPC controllers can only help reduce the central plant costs (the campus costs are exogenous).

From Table 4.2 we also see that the expected cost of the central plant is improved by 7.52% by using stochastic MPC (relative to deterministic MPC). The associated cost savings total 349,000 $/year. These savings represent 75% of the possible improvement over deterministic MPC (obtained with perfect information MPC).

Table 4.3: Economic performance analysis for different MPC implementations (central plant only).

	Campus Only (No Central Plant)	**Central Plant Only**			
		Perfect Information	Stochastic $\beta = 0$	Deterministic $\beta = 0.1$	Value of Stochastic MPC
Expected Electricity Cost ($/year)	9,764,251	1,605,815	1,658,335	1,780,771	122,436 **(6.88%)**
Expected Water Cost ($/year)	-	1,780,554	1,796,342	1,813,681	17,339 **(0.96%)**
Expected Natural Gas Cost ($/year)	-	523,210	545,679	596,801	51,122 **(8.57%)**
Expected Demand Charge ($/year)	2,051,316	222,015	315,637	449,206	133,569 **(29.8%)**
Cost of Central Plant ($/year)	-	4,231,595	4,291,497	4,640,460	348,963 **(7.52%)**
Savings in Central Plant Cost	-	9.66%	7.52%	**Base**	7.52%

Table 4.3 disaggregates the costs of the central plant. Here, we observe that stochastic MPC achieves an improvement of 29.8% in the demand charge cost over deterministic MPC and achieves improvements in electricity cost of 6.88%, in natural gas cost of 8.57%,

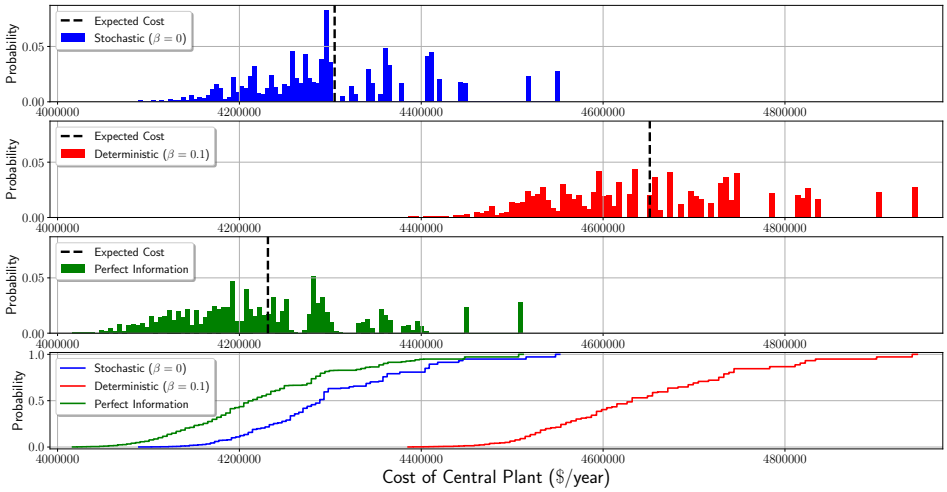

Figure 4.12: Probability and cumulative distributions for total cost of central plant.

and in water cost of 0.96%. We thus see that, even if reductions in electricity use are moderate, reductions in cost are significant. We also note that reductions in natural gas use and cost are both 8.57% (because the price of gas is not varying over time). The dramatic reduction in demand charge costs indicates that disturbance uncertainty has a strong effect on peak electricity load (we recall that disturbance uncertainty rises sharply over time).

The probability distribution and cumulative distribution for the costs of the HVAC central plant and the peak demand charges are shown in Figures 4.12 and 4.13. From the cumulative distributions, we observe that the probability of taking a smaller cost and a smaller demand charge for the central plant is higher under stochastic MPC compared to deterministic MPC. Similarly, the probability of obtaining a higher cost and a higher demand charge for the central plant is much higher under deterministic MPC than under stochastic MPC. This clearly illustrates that the performance of stochastic MPC consistently dominates that of deterministic MPC.

We then evaluated constraint violations (in terms of storage overflow or drying up) obtained with deterministic and stochastic MPC for all the validation scenarios. We recall that the MPC controllers take no action when the control policy becomes infeasible. This

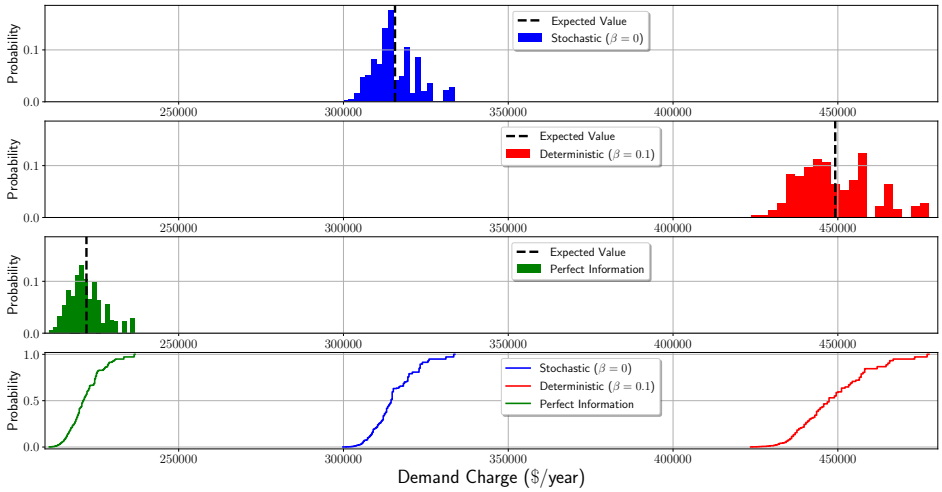

Figure 4.13: Probability and cumulative distributions for demand charges.

can result in either overflow or drying up of the chilled water and hot water storage tanks, and therefore can lead to loss of energy. Figure 4.14 shows how often infeasibility occurs per 100 hours over the year for the 200 validation scenarios. The results highlight that stochastic MPC without a storage buffer is more reliable at maintaining a feasible operation than deterministic MPC with a 10% buffer.

In the above case study, we chose a buffer of 10% for deterministic MPC as this was the best possible buffer value found. To see this, we examined the effect of using different buffers on economic performance. Table 4.4 compares results for buffers of 0-20% for deterministic MPC. As can be seen, deterministic MPC performs poorly without any buffer and results in a very high cost (compared to performance obtained with a buffer of 10%). This is because a deterministic MPC implementation with 0% storage buffer tries to utilize the full capacity of the TES in each optimization step but, after the actual realization of the loads is observed, it has to adjust the control actions because of frequent constraint violations which in turn leads to higher demand charges. Increasing the buffer initially leads to lower total cost for deterministic MPC because the controller is able to avoid infeasibility issues but eventually becomes detrimental because the fraction of the storage capacity

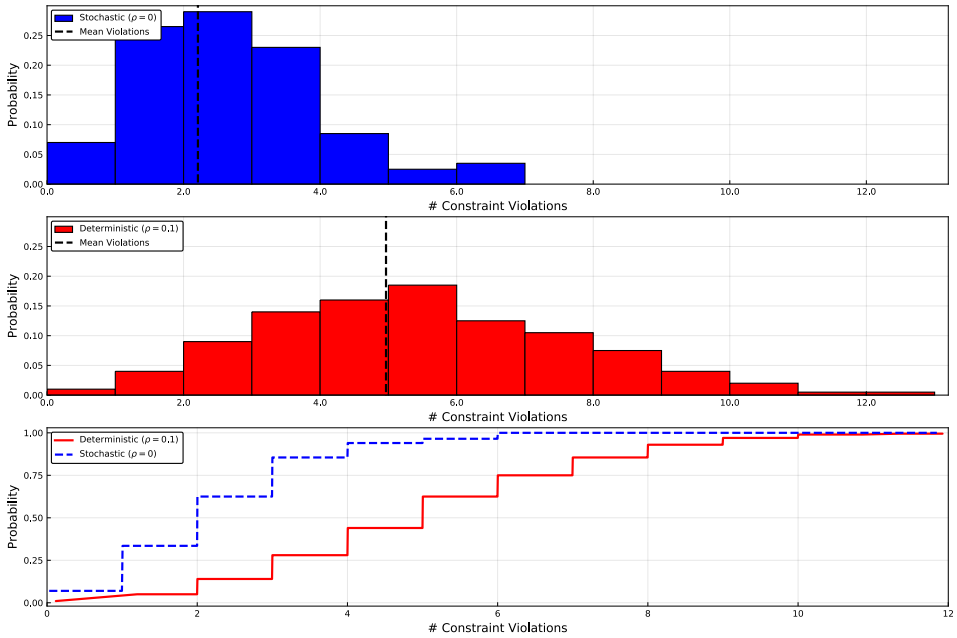

Figure 4.14: Probability and cumulative distributions for constraint violations (per 100 hours).

available decreases (thus limiting flexibility). This inherent trade-off is shown in Table 4.4 and Figures 4.15. Here, it is clear that a buffer of 10% achieves the best cost. We note, however, that even this best tuned cannot achieve the performance of stochastic MPC. In other words, stochastic MPC consistently dominates deterministic MPC.

Table 4.4: Expected costs for deterministic MPC with varying buffers.

Item	**Deterministic** ($\rho = 0$)	**Deterministic** ($\beta = 0.05$)	**Deterministic** ($\beta = 0.08$)	**Deterministic** ($\beta = 0.1$)	**Deterministic** ($\beta = 0.13$)	**Deterministic** ($\beta = 0.15$)	**Deterministic** ($\beta = 0.2$)
Total Cost ($/year)	17,583,658	17,050,005	16,558,117	16,456,027	16,604,521	17,212,421	17,950,255
Cost of Central Plant ($/year)	5,768,091	5,234,438	4,742,550	4,640,461	4,788,954	5,396,854	6,134,688
Improvement in Cost of Central Plant	-24.3%	-12.8%	-2.20%	**Base**	-3.21%	-16.3%	-32.2%

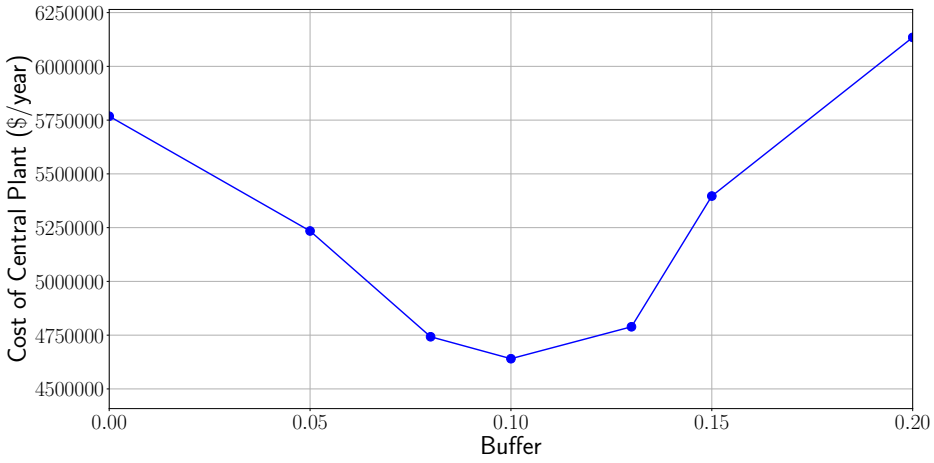

Figure 4.15: Economic performance (cost of central plant) of deterministic MPC with varying buffers.

4.4 Conclusions

In this chapter, we presented a stochastic MPC framework for a central HVAC plant of a typical university campus. We use the framework to rigorously assess the benefits of stochastic MPC over deterministic MPC in terms of economic performance and constraint violations. Our framework uses real historical data to conduct forecasting and uncertainty quantification tasks. Our results show that stochastic MPC reduces overall energy consumption and cost by better handling of storage and better integration of electricity, natural gas, and water. Specifically, we found that stochastic MPC can reduce the natural gas cost by 8.57%, electricity cost by 6.89%, and peak demand charges by 29.8% (relative to deterministic MPC). We find that deterministic MPC leads to frequent constraint violations in storage capacity (causing overflow or drying up) of the hot and chilled water tanks. Stochastic MPC is able to avoid constraint violations because it anticipates the uncertainty by explicitly incorporating it in the model.

5

HIERARCHICAL MPC SCHEMES FOR PERIODIC SYSTEMS USING STOCHASTIC PROGRAMMING

In this chapter, we design a hierarchical model predictive control (MPC) framework using stochastic programming to handle long horizon (or infinite horizon) problems for periodic systems. This is based on the observation that if the state policy of an infinite-horizon problem is periodic, the problem can be cast as a stochastic program (SP). We show that it is possible to update periodic state targets by solving a retroactive optimization problem that progressively accumulates historical data. Moreover, we show that the retroactive problem is a statistical approximation of the SP and thus delivers optimal targets in the long run. Notably, the computation of the optimal targets can be achieved without data forecasts. The SP setting also reveals that the retroactive problem can be seen as a high-level hierarchical layer that provides targets to guide a low-level MPC controller that operates over a short period at high time resolution. We derive a retroactive scheme tailored to linear systems by using cutting plane techniques and suggest strategies to handle nonlinear systems and analyze stability properties (Kumar et al., 2019b).

5.1 Introduction

A well-known challenge arising in model predictive control (MPC) is the computational complexity associated with the length of the planning horizon and with the time resolution of the state and control policies (Rawlings et al., 2017). These issues are often encountered in energy system applications that exhibit phenomena and disturbances emanating at multiple timescales. For instance, in energy systems, long horizons are often required to respond to low-frequency (e.g., seasonal) supply/demand variations and peak electricity costs (e.g., demand charges) while fine time resolutions are needed to modulate high-frequency variations (e.g., from wind/solar supply) and to participate in real-time markets (Dowling et al., 2017; Braun, 1990). Computational complexity issues are often handled using receding horizon (RH) approximations, which are practical but do not provide optimality guarantees (Risbeck et al., 2017).

Hierarchical MPC schemes (Scattolini and Colaneri, 2007; Picasso et al., 2010) have been recently proposed to handle multiple scales (and associated computational complexity issues). These schemes, however, do not provide optimality guarantees in the sense that the computed policies match those of the long-horizon problem of interest. The hierarchical scheme proposed in Zavala (2016) uses adjoint information obtained from a long-term but coarse controller to guide a short-term controller operating at fine time resolutions. Computational experiments are provided to demonstrate that this approach can achieve optimality but no guarantees are given. Moreover, such an approach requires smoothness and continuity of the adjoint profiles, which is not guaranteed in general applications.

The hierarchical scheme proposed in this work relies on the observation that, if the optimal policy of an infinite horizon problem is periodic (or can be approximated with a periodic policy), the problem can be cast as a stochastic programming (SP) problem. Periodicity is a property that is commonly observed in systems driven by exogenous factors with strong periodic components (e.g., energy demands and prices) (Huang et al., 2011; Risbeck et al., 2015). Under the SP abstraction, the inter-period trajectory of the exogenous

factors is interpreted as a realization of a random variable that triggers a periodic trajectory of the system states (the states at the beginning and end of the period are the same). Moreover, the periodic states are interpreted as design variables, and operational policies over the periods are interpreted as recourse variables. We have recently observed that the SP representation provides a mechanism to construct hierarchical MPC schemes in which a *long-term* (supervisory) MPC controller provides periodic targets to guide a *s*hort-term MPC controller (Kumar et al., 2018d,b) (see Appendix Band Appendix C). In Section B.4 of Appendix B, we show that under nominal conditions with *perfect forecasts*, the hierarchical scheme delivers an optimal policy for the infinite horizon problem. For the more relevant case of *imperfect forecasts*, the hierarchical scheme needs to re-compute periodic targets. While this can certainly be done using an RH scheme (e.g., computes targets by anticipating multiple future periods), such an approach would not provide optimality guarantees. In fact, to the best of our knowledge, no RH scheme currently exists that can provide optimal policies in the presence of imperfect forecast information. Specifically, standard proactive RH schemes use historical data to compute forecasts and associated control actions. A fundamental issue with proactive approaches is that no optimality guarantees can be provided unless the forecast is perfect.

The contribution of this work is the observation that, under a periodic setting, one can derive *retroactive* hierarchical MPC schemes that accumulate real historical data to *asymptotically deliver optimal targets*. The retroactive design principle thus offers optimality guarantees and, notably, does not require data forecasts. The retroactive approach thus provides key advantages over proactive RH schemes. The targets obtained with the retroactive scheme are used to guide a low-level controller operating at fine time resolutions within the periods. In the case of linear systems, one can derive a specialized retroactive scheme by using *incremental cutting-plane (CP) techniques* (Higle and Sen, 1991). The SP setting also reveals strategies to construct retroactive schemes for nonlinear systems and to obtain the desired stability properties. We demonstrate the concepts using a battery application and compare the proposed retroactive hierarchical MPC scheme with a proactive

MPC approach for periodic systems.

The chapter is structured as follows. In Section 5.2, we provide basic definitions and describe the problem setting. In Section 5.3, we introduce the concept of retroactive optimization, derive an incremental CP scheme for linear systems, discuss implementation details, and discuss extensions for nonlinear systems. Computational experiments are presented in Section 5.4.

5.2 Basic Definitions and Setting

In this work, we derive schemes to compute approximation solutions for the long-horizon problem $\mathbf{O}_m$:

$$\min_{u_{\xi}, x_{\xi}, \eta, x_0} \ \frac{1}{m} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) + \eta$$

$$\text{s.t. } \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \leq \eta, \ \xi \in \Xi, t \in \mathcal{T} \tag{5.2.1a}$$

$$x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}), \ \xi \in \Xi, t \in \bar{\mathcal{T}} \tag{5.2.1b}$$

$$x_{\xi+1,0} = x_{\xi,n}, \ \xi \in \bar{\Xi} \tag{5.2.1c}$$

$$x_{1,0} = x_0 \tag{5.2.1d}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{5.2.1e}$$

Here, the horizon with $p := (n+1) \cdot m$ time steps is partitioned into a set of *periods* $\Xi := \{1, \ldots, m\}$ with intra-period times $\mathcal{T} := \{0, \ldots, n\}$. For convenience, we define the sets $\bar{\mathcal{T}} := \mathcal{T} \setminus \{n\}$ and $\bar{\Xi} := \Xi \setminus \{m\}$. The controls, states, and data trajectories at period $\xi \in \Xi$ and intra-period time $t \in \mathcal{T}$ are denoted as $u_{\xi,t} \in \mathbb{R}^{n_u}$, $x_{\xi,t} \in \mathbb{R}^{n_x}$, and $d_{\xi,t} \in \mathbb{R}^{n_d}$, respectively. We define the notation $u_\xi := \{u_{\xi,t}\}_{t \in \mathcal{T}}$, $x_\xi := \{x_{\xi,t}\}_{t \in \mathcal{T}}$, and $d_\xi := \{d_{\xi,t}\}_{t \in \mathcal{T}}$ to denote the inner period trajectories. The problem *data* is interpreted as exogenous disturbances or factors driving the system (e.g., market prices, demands, weather, model errors).

The mapping $\varphi_1 : \mathbb{R}^{n_u} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_d} \to \mathbb{R}$ is a time-additive cost function and the mapping $\varphi_2 : \mathbb{R}^{n_u} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_d} \to \mathbb{R}$ is a time-max (peak) cost function. We assume both of these functions to be bounded in their domains. Minimizing the variable $\eta \in \mathbb{R}$ subject to the constraints (5.2.1a) is equivalent to minimize the peak cost $\max_{\xi \in \Xi} \max_{t \in \mathcal{T}} \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t})$. Consequently, since $\varphi_2(\cdot)$ is assumed to be bounded, we have that η is bounded as well. The mapping $f : \mathbb{R}^{n_u} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_d} \to \mathbb{R}^{n_x}$ describes the system dynamics and $\mathcal{X}$ and $\mathcal{U}$ are non-empty feasible sets for states and controls, respectively. For reasons that will become apparent, the initial state $x_0 \in \mathbb{R}^{n_x}$ in the above formulation is treated as a decision variable.

The infinite-horizon problem $\mathbf{O}_\infty$ is obtained by setting $\lim_{m \to \infty}$. If the data $d_\xi, \xi \in \Xi$ is known, the solution of $\mathbf{O}_\infty$ provides a policy with the *best possible performance*. Unfortunately, problem $\mathbf{O}_m$ becomes computationally difficult to solve for large m (long horizons) and/or for large n (fine time resolutions). An approximate solution for $\mathbf{O}_\infty$ is often computed by using a *proactive RH scheme*, which approximates the policy by moving forward in time and by planning over short time horizons (e.g., that span a few periods). Our goal is to derive an *alternative* strategy that uses hierarchical MPC schemes that approximate the solution of $\mathbf{O}_\infty$ by using a *periodic policy*. We define a periodic policy as follows:

Definition 5.2.1. *A policy is said to be periodic if it satisfies the periodicity constraints $x_{\xi,0} = x_{\xi,n}$ for all $\xi \in \Xi$.*

To obtain a periodic policy, we enforce periodicity constraints at every period (after every n steps). These constraints, together with the continuity constraints (5.2.1c), can be expressed as $x_{\xi+1,0} = x_{\xi,0}$, $\xi \in \bar{\Xi}$. This set of constraints, in turn, can be reformulated as $x_{\xi,0} = x_0$, $\xi \in \Xi$. These modifications give the periodic problem $\mathbf{P}_m$:

$$\min_{u_\xi, x_\xi, \eta, x_0} \ \frac{1}{m} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) + \eta \tag{5.2.2a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \leq \eta, \ \xi \in \Xi, t \in \mathcal{T} \tag{5.2.2b}$$

$$x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}), \ \xi \in \Xi, t \in \bar{\mathcal{T}} \tag{5.2.2c}$$

$$x_{\xi,0} = x_0,\ \xi \in \Xi \tag{5.2.2d}$$

$$x_{\xi,n} = x_0,\ \xi \in \Xi \tag{5.2.2e}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{5.2.2f}$$

Remark: The feasible region of $\mathbf{P}_m$ is smaller than that of $\mathbf{O}_m$ (since the latter does not enforce periodicity every n steps). Consequently, the performance of $\mathbf{P}_m$ is expected to be inferior to that of $\mathbf{O}_m$. In some applications, however, the deterioration of performance might not be significant. For instance, in Kumar et al. (2018d), it is shown that enforcing state periodicity constraints for a battery system (obtained from $\mathbf{P}_m$) results in a policy with a cost that is 0.2% larger than the cost achieved without periodicity constraints (obtained from $\mathbf{O}_m$).

The solution of $\mathbf{P}_m$ provides a better approximation to the solution of $\mathbf{O}_m$ as we increase n (increasing the size of the period). This is because periodicity constraints are enforced less often. For a fixed value horizon p, we have that in the limit when $n = p$ and $m = 1$, we obtain the best possible periodic policy (enforcing periodicity at the beginning and the end of the horizon). On the other hand, in the limit when $n = 1$ and $m = p$, we have the worst possible periodic policy (a steady-state policy). Setting $n = p$ and $m = 1$ and eliminating all the periodicity constraints make $\mathbf{P}_m$ and $\mathbf{O}_m$ equivalent. From these observations, we note that the length of the period n can be used as a design parameter to find a periodic policy of $\mathbf{P}_m$ that properly approximates the policy of $\mathbf{O}_m$.

The goal of $\mathbf{P}_m$ is to find a periodic state x_0, peak cost η, as well as intra-period policies $u_\xi, x_\xi\ \xi \in \Xi$ that minimize the time-additive and peak costs. We note that, by construction, $x_{\xi+1,0} = x_{\xi,0} = x_0$ holds for all $\xi \in \Xi$ at any solution of $\mathbf{P}_m$. If we have perfect knowledge of the data $d_\xi, \xi \in \Xi$, the infinite-horizon problem $\mathbf{P}_\infty$ identifies a periodic policy with the best possible performance. We denote a solution of $\mathbf{P}_\infty$ as $w^* = (x_0^*, \eta^*)$, and u_ξ^*, x_ξ^*. Unfortunately, problem $\mathbf{P}_m$ also becomes computationally difficult to solve for large m and/or n. Here, one can address this issue by using a proactive RH scheme with periodicity constraints, as proposed in Huang et al. (2011). We will see, however, that periodicity results

in a structure that enables the derivation of *retroactive schemes* that offer key advantages over proactive schemes. In particular, we will see that retroactive schemes can deliver a solution of $\mathbf{O}_\infty$ (while proactive schemes cannot).

By analyzing the structure of $\mathbf{P}_m$, we can see that the only coupling between periods arises from the variables x_0 and η. Consequently, by fixing these variables, we can decompose (5.2.2) into individual *period subproblems* of the form:

$$\min_{u_\xi, x_\xi} \sum_{t \in \mathcal{T}} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \tag{5.2.3a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \leq \eta, \ t \in \mathcal{T} \tag{5.2.3b}$$

$$x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}), \ t \in \bar{\mathcal{T}} \tag{5.2.3c}$$

$$x_{\xi,0} = x_0 \tag{5.2.3d}$$

$$x_{\xi,n} = x_0 \tag{5.2.3e}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{5.2.3f}$$

We define this problem as $\mathbf{S}_\xi$ and we define its optimal cost as $h(w, d_\xi)$, where $w := (x_0, \eta) \in \mathcal{W} := \mathcal{X} \times \mathbb{R}$. In the following discussion, we use the short-hand notation $h_\xi := h(w, d_\xi)$.

We now make the following key assumption regarding the nature of the data and of the subproblem cost $h(w, d_\xi)$.

Assumption 5.2.1. *Assume that $\{d_\xi\}_{\xi=1}^\infty$ is a sequence of independent and identically distributed (i.i.d.) realizations of a continuous random variable D with associated probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with associated codomain $\Omega' \subseteq \mathbb{R}^{n_d}$ and σ-algebra $\mathcal{F}'$. Moreover, assume that the cost function $h : \mathcal{W} \times \Omega' \to \mathbb{R}$ is continuous and bounded in its domain.*

Under the assumption of i.i.d. realizations and a bounded function $h(\cdot)$, we can establish that the sample average $\frac{1}{m}\sum_{\xi=1}^m h(w, d_\xi)$ converges pointwise with probability one to $\mathbb{E}[h(w, D)]$ as $m \to \infty$. Here, $\mathbb{E}[\cdot]$ is the expectation operator. This property is the strong law of large numbers (LLN) and is key because it reveals that the infinite horizon problem

$\mathbf{P}_\infty$ can be interpreted as a *stochastic programming (SP) problem* of the form:

$$\min_{w \in \mathcal{W}} \phi(w) := g(w) + \mathbb{E}\left[h(w, D)\right]. \tag{5.2.4}$$

Under this representation, periods are realizations d_ξ of D, $w = (x_0, \eta)$ are *design* (target) variables, and x_ξ, u_ξ are (recourse) policies associated with realization d_ξ. Here, $g : \mathcal{W} \to \mathbb{R}$ is a cost function for the design variables. In the context of $\mathbf{P}_\infty$, we have $g(w) = \eta$, but this function can be generalized to enforce a cost also on x_0. Function $g(\cdot)$ is bounded because $\phi_2(\cdot)$ (and thus η) are assumed to be bounded. Since (5.2.4) is equivalent to $\mathbf{P}_\infty$, it delivers optimal targets w^*. We define the solution set of (5.2.4) as $S \subseteq \mathcal{W}$.

Remark: The i.i.d. requirement on $\{d_\xi\}$ ensures that LLN holds. This requirement can be enforced by defining a sufficiently long period duration n (to eliminate autocorrelation between periods). This approach is commonly used in statistical extrapolation and time series analysis (Box et al., 2015; Ragan and Manuel, 2008). The independence requirement can also be relaxed by allowing for cost sequences that have bounded correlation (Hu et al., 2008). To give an idea of why this is the case, we provide the following result.

Property 5.2.1. *Assume that* $\{h_\xi\}_{-\infty}^{\infty}$ *is a sequence of identically distributed random variables with expected value* $\mu = \mathbb{E}[h(x, D)]$. *Assume also that there exists* $0 \leq c < \infty$ *such that* $\sum_{\xi=-\infty}^{\infty} \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \leq c$ *holds for all* $\xi' = 1, \ldots, \infty$. *Then,* $\lim_{m\to\infty} \frac{1}{m} \sum_{\xi=1}^{m} h_\xi = \mathbb{E}[h(x, D)]$.

Proof: Define $S_m := \sum_{\xi=1}^{m} h_\xi$, $\bar{h}_m := S_m/m$, and $\mu := \mathbb{E}[h(x, D)]$. The variance of S_m (denoted as $V[S_m]$) is:

$$\begin{aligned}
\mathbb{E}[(S_m - m\,\mu)^2] &= \sum_{\xi'=1}^{m} \sum_{\xi=1}^{m} \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \\
&\leq \sum_{\xi'=1}^{m} \sum_{\xi=-\infty}^{\infty} \mathbb{E}[(h_{\xi'} - \mu)(h_\xi - \mu)] \\
&\leq m\,c.
\end{aligned}$$

We thus have $V[\bar{h}_m] \leq c/m$ and, from Chebyshev's inequality (with parameter $\kappa > 0$),

$$\begin{aligned}\mathbb{P}(|\bar{h}_m - \mu| > \kappa) &\leq \frac{1}{\kappa^2}\mathbb{E}[(\bar{h}_m - \mu)^2] \\ &= \frac{1}{\kappa^2 m^2}\mathbb{E}[(S_m - m\mu)^2] \\ &\leq \frac{m\,c}{m^2\kappa^2} = \frac{c}{m\,\kappa^2}.\end{aligned}$$

We thus have $\lim_{m\to\infty} \bar{h}_m = \mu$ with probability one. □

Remark: We can guarantee that a solution for $\mathbf{P}_m$ exists if the period length n is long enough such that one can find a control policy u_ξ that delivers a feasible solution for $\mathbf{S}_\xi$ for all fixed $w \in \mathcal{W}$ and $d_\xi, \xi \in \Xi$. Here, a feasible solution of $\mathbf{S}_\xi$ is one that satisfies the control and state constraints as well as the periodicity and peak constraints. This assumption is compatible with that used in the standard MPC literature (a sufficiently long horizon is assumed such that one can find a feasible control policy that satisfies the state and terminal constraints) (Subramanian et al., 2014; Zanon et al., 2017). In the SP literature, such an assumption is equivalent to assuming that $\mathbf{S}_\xi$ has *relatively complete recourse* (i.e., $\mathbf{S}_\xi$ has a feasible solution for any $w \in \mathcal{W}$ and $d_\xi \in \Omega'$).

5.3 Hierarchical MPC Schemes

The SP representation opens a number of interesting directions. In particular, it provides a mechanism to derive *hierarchical MPC schemes*. For instance, one can use a statistical approximation of (5.2.4) (equivalently of $\mathbf{P}_\infty$) to provide targets x_0, η that guide a short-term MPC controller of the form (5.2.3). As we discuss next, approximations of $\mathbf{P}_\infty$ can be constructed and solved in a tractable manner by using well-established SP techniques (Zavala et al., 2008; CarøE and Schultz, 1999; Geoffrion, 1972).

5.3.1 *Retroactive Optimization*

A key observation that arises from the SP representation is that one can derive *retroactive optimization* schemes that accumulate data over time to refine targets. To explain how such a scheme would work, assume that the system is currently at the beginning of period $m+1$ and that the data history $\{d_\xi\}_{\xi=1}^m$ is known. We use this information to solve the problem:

$$\min_{w \in \mathcal{W}} \ \phi_m(w) := g(w) + \frac{1}{m} \sum_{\xi=1}^{m} h(w, d_\xi). \tag{5.3.5}$$

This problem is equivalent to $\mathbf{P}_m$, and because $\{d_\xi\}_{\xi=1}^m$ is i.i.d., $\mathbf{P}_m$ is a statistical approximation of $\mathbf{P}_\infty$. In the SP literature, problem (5.3.5) is known as a sample average approximation (SAA). We define the solution set of $\mathbf{P}_m$ as $S_m \subseteq \mathcal{W}$. A solution of $\mathbf{P}_m$ is used to update the targets for the next period $w_{m+1} = (x_{m+1,0}, \eta_{m+1})$. A solution of $\mathbf{P}_m$ also implicitly contains optimal (retroactive) policies u_ξ, x_ξ, $\xi = 1, \ldots, m$ associated with the historical sequence $\{d_\xi\}_{\xi=1}^m$. These retroactive policies are interpreted as policies that the system *would have taken* given knowledge of the data.

Given that the system is at the current state target $x_{m,0}$ (obtained from the solution of $\mathbf{P}_{m-1}$), we use the targets w_{m+1} to guide a *short-term MPC controller* over period $m+1$. At this point, however, the data d_{m+1} is not known, so we use a *forecast* $\hat{d}_{m+1}$ to find policies that optimize the system over the next period $m+1$ while satisfying the targets w_{m+1}. This can be interpreted as solving $\mathbf{S}_\xi$ for w_{m+1} given $\hat{d}_{m+1}$. The forecast $\hat{d}_{m+1}$ is typically obtained by using forecasting techniques such as AR, ARMA, or ARIMA time series models, or covariance estimators (Kumar et al., 2018a). At the beginning of the next period $m+2$, the actual data realization d_{m+1} reveals itself and we use this to solve the approximation $\mathbf{P}_{m+1}$ to obtain new targets w_{m+2}.

The retroactive scheme is consistent because, from the law of large numbers, we know that accumulating data over time will yield an asymptotically exact statistical approximation $\lim_{m\to\infty} \phi_m(w) = \phi(w)$ and thus the targets obtained with $\mathbf{P}_m$ will provide a solution

to $\mathbf{P}_\infty$ as $m \to \infty$. This asymptotic convergence result is formally stated in the following theorem (see Theorem 5.3 in Shapiro et al. (2009)).

Theorem 5.3.1. *Suppose that there exists a compact set $C \subset \mathbb{R}^{n_w}$ such that: i) the solution set S of $\mathbf{P}_\infty$ is nonempty and contained in C, ii) the function $\phi(w)$ is finite-valued and continuous on C, iii) $\phi_m(w)$ converges to $\phi(w)$ with probability one as $m \to \infty$ uniformly in $w \in C$, and with probability one for large enough m the solution set S_m of $\mathbf{P}_m$ is nonempty and contained in C. Then, the Hausdorff distance between the solution sets $\mathbb{D}(S_m, S)$ converges to zero with probability one as $m \to \infty$.*

The above result implies that the distance of any solution of $\mathbf{P}_m$ to the solution set of $\mathbf{P}_\infty$ converges to zero as $m \to \infty$. Statistical approximation results for SPs also indicate that the probability that a solution of $\mathbf{P}_m$ is in the solution set of $\mathbf{P}_\infty$ increases exponentially with m (Theorem 5.16 in Shapiro et al. (2009)). In other words, the *probability of finding better targets than those obtained with $\mathbf{P}_m$ decays exponentially as information is accumulated over time*.

These asymptotic optimality results provide a key advantage of the retroactive scheme over traditional proactive RH schemes. This is based on a fundamental design difference: the retroactive scheme uses past (but real) data while proactive RH schemes use future (but approximate) data forecasts. Moreover, proactive schemes discard historical data when computing new targets. The fact that historical data is discarded prevents RH schemes from offering asymptotic optimality guarantees.

5.3.2 *Incremental Cutting Plane Scheme*

The structure of $\mathbf{P}_m$ can be exploited using decomposition strategies and this enables scalability to large values of m. In this section, we provide a decomposition scheme for linear systems, and we then discuss potential extensions to nonlinear systems.

The retroactive scheme for linear systems proposed is based on an *incremental cutting plane* (CP) scheme. Our approach is an adaptation of the stochastic decomposition scheme proposed in Higle and Sen (1991) to tackle linear SPs. To derive this linear setting, we

assume that $g(w) = c_w^T w$ where $c_w \in \mathbb{R}^{n_w}$ is a cost vector, we assume that the set $\mathcal{W} \subseteq \mathbb{R}^{n_w}$ is polyhedral, and we assume that $\mathbf{S}_\xi$ has the form:

$$h(w, d_\xi) := \min_{y \in \mathbb{R}_+^{n_y}} c_\xi^T y \quad \text{s.t.} \quad Wy = r_\xi - Tw. \tag{5.3.6}$$

Here, the data realization is given by $d_\xi = (c_\xi, r_\xi)$. We use $y(w, d_\xi)$ to denote the primal solution vector containing the intra-period trajectories (x_ξ, u_ξ) and some additional dummy variables. The intra-period dynamics are captured using W, T that are coefficient matrices. The structure of the recourse problem is used to simplify algebraic manipulations and is done without loss of generality (the results that we present hold provided that the recourse problem is a linear program). The representation of $\mathbf{S}_\xi$ allows us to express its dual form in the following compact form:

$$\max_{\pi} \ \pi^T (r_\xi - Tw) \quad \text{s.t.} \quad W^T \pi \leq c_\xi. \tag{5.3.7}$$

Here, $\pi(w, d_\xi)$ is a dual solution vector (dual variables of $\mathbf{S}_\xi$) and we recall that $\pi(w, d_\xi)^T (r_\xi - Tw) = h(w, d_\xi)$ holds for $w \in \mathcal{W}$ (from strong duality). We assume that the feasible set of the dual subproblem is a non-empty, compact, and convex polyhedral set, and therefore the polyhedron represented by the set $\mathcal{P} := \{\pi \mid W^T \pi \leq c_\xi\}$ is a pointed polyhedron for all $(w, d_\xi) \in \mathcal{W} \times \Omega'$. As a result, $\mathbf{S}_\xi$ has a finite number of dual vertices or extreme points (where a vertex of the polyhedron $\mathcal{P}$ is a vector $\pi \in \mathcal{P}$ such that we cannot find two distinct vectors $\pi_1, \pi_2 \in \mathcal{P}$ and a scalar $\lambda \in [0, 1]$, such that $\pi = \lambda \pi_1 + (1 - \lambda)\pi_2$ (see Chapter 2 of Bertsimas and Tsitsiklis (1997)). Moreover, because the support Ω is finite, the set of dual vertices for all subproblems $\mathbf{S}_\xi$ is finite (see Theorem 2.9 in Chapter 2 of Bertsimas and Tsitsiklis (1997)). We use $\mathbb{V}$ to denote the set of all these dual vertices. Consequently, by definition, $\pi(w, d_\xi) \in \mathbb{V}$ for all $(w, d_\xi) \in \mathcal{W} \times \Omega'$.

In the linear setting, the cost function of $\mathbf{P}_\infty$ (given by $\phi(w) = c_w^T w + \mathbb{E}[h(w, D)]$) can

be outer-approximated using CPs accumulated over $\xi = 1, \ldots, m$ as:

$$\underline{\phi}_m(w) := \max\{\alpha_\xi^m + (c_w + \beta_\xi^m)^T w | \xi = 1, \ldots, m\}, \tag{5.3.8}$$

where the coefficients α_ξ^m, β_ξ^m are selected to match:

$$\alpha_\xi^m + (\beta_\xi^m)^T w = \frac{1}{m} \sum_{\xi=1}^{m} (\pi_\xi^m)^T (r_\xi - Tw). \tag{5.3.9}$$

Here, $\pi_\xi^m \in \text{argmax}\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}_m\}$ for $\xi = 1, \ldots, m$, where $\mathbb{V}_m \subseteq \mathbb{V}$ is the collection of vertices accumulated up to period m and $w_m = (x_{m,0}, \eta_m)$. For convenience, we define the function $h_m(w, d_\xi) := \max\{\pi^T(r_\xi - Tw) | \pi \in \mathbb{V}_m\}$ and note that $h_m(w, d_\xi) = (\pi_\xi^m)^T(r_\xi - Tw)$ holds. Moreover, we note that $h(w, d_\xi) = \max\{\pi^T(r_\xi - Tw) | \pi \in \mathbb{V}\}$ holds. Since $\mathbb{V}_m \subseteq \mathbb{V}$, we have that $(\pi_\xi^m)^T(r_\xi - Tw) \leq h(w, d_\xi)$.

The *running cost* is given by $\phi_m(w) = c_w^T w + \frac{1}{m} \sum_{\xi=1}^{m} h(w, d_\xi)$. We will prove that $\underline{\phi}_m(w)$ underestimates the running cost $\phi_m(w)$ for all m and converges to the infinite-horizon cost $\phi(w)$ as $m \to \infty$. Consequently, at period m, we update the targets by solving the *master* problem $\mathbf{M}_m$:

$$\min_{w \in \mathcal{W}} \ \underline{\phi}_m(w). \tag{5.3.10}$$

This problem is a tractable surrogate of $\min_{w \in \mathcal{W}} \ \phi_m(w)$, because it captures the recourse subproblems by using hyperplanes. This becomes particularly important as information is accumulated over time. The solution of $\mathbf{M}_m$ is used to update the targets w_{m+1}, which in turn are used to solve the recourse subproblem $S_{\xi+1}$ and with this, obtain a new dual vertex to be stored in $\mathbb{V}_{m+1}$. The CP scheme is summarized as:

1. Initialize $m \leftarrow 1$, $\mathbb{V}_m \leftarrow \emptyset$, and w_m.

2. At period time $m + 1$:

3. Observe d_m and solve $\mathbf{S_m}$ to obtain $\pi(w_m, d_m)$.

4. Update $\mathbb{V}_m \leftarrow \mathbb{V}_{m-1} \cup \{\pi(w_m, d_m)\}$.

5. (a) Obtain $\pi_\xi^m \in \text{argmax}\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}_m\}$ for all $\xi = 1, \ldots, m$.

   (b) Get α_m^m and β_m^m from (5.3.9).

   (c) Update $\alpha_\xi^m \leftarrow \frac{m-1}{m}\alpha_\xi^{m-1}$, $\beta_\xi^m \leftarrow \frac{m-1}{m}\beta_\xi^{m-1}$ for $\xi = 1, \ldots, m-1$.

6. Solve $\mathbf{M}_m$ and obtain updated targets w_{m+1}.

7. Shift period time $m \leftarrow m+1$ and return to Step 2.

We now prove that the CP scheme delivers a sequence of targets $\{w_m\}_{m=1}^\infty$ that converges to optimal targets w^* of $\mathbf{P}_\infty$. Our analysis follows along the lines of that presented in Higle and Sen (1991).

Theorem 5.3.2. *The CP $\alpha_m^m + (c_w + \beta_m^m)^T w_m$ generated in period m provides a statistically valid lower bound (statistically based estimate of a lower bound) for $\phi(w)$ for all $w \in \mathcal{W}$.*

Proof: Because $\mathbb{V}_m \subseteq \mathbb{V}$ we have that,

$$\max\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}_m\} \leq \max\{\pi^T(r_\xi - Tw_m) | \pi \in \mathbb{V}\},$$

and $(\pi_\xi^m)^T(r_\xi - Tw_m) = h(w_m, d_\xi)$, $\xi = 1, \ldots, m$. We thus have

$$\frac{1}{m}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw_m) \leq \frac{1}{m}\sum_{\xi=1}^m h(w_m, d_\xi).$$

Furthermore, $\pi^T(r_\xi - Tw) \leq h(w, d_\xi)$ for any $\pi \in \mathbb{V}$ and:

$$c_w^T w + \frac{1}{m}\sum_{\xi=1}^m (\pi_\xi^m)^T(r_\xi - Tw) \leq c_w^T w + \frac{1}{m}\sum_{\xi=1}^m h(w, d_\xi) = \phi_m(w),\ w \in \mathcal{W}.$$

The result follows from (5.3.9) and by noticing that $\phi_m(w)$ is a statistical approximation of $\phi(w)$. □

As more observations d_ξ are collected, it is important that all the collected CPs provide a statistically valid lower bound for $\phi(w)$. This is the goal of Step 5c in the CP scheme. In

particular, at period $m+p$ with $p>0$:

$$\begin{aligned}\frac{1}{m+p}\sum_{\xi=1}^{m}(\pi_{\xi}^{m})^{T}(r_{\xi}-Tw) &\leq \frac{1}{m+p}\sum_{\xi=1}^{m}h(w,d_{\xi}) \\ &\leq \frac{1}{m+p}\sum_{\xi=1}^{m+p}h(w,d_{\xi})\end{aligned}$$

Thus, in the $(m+p)^{th}$ period, the CP

$$c_{w}^{T}w+\frac{1}{m+p}\sum_{\xi=1}^{m}(\pi_{\xi}^{m})^{T}(r_{\xi}-Tw)=\alpha_{m}^{(m+p)}+(\beta_{m}^{(m+p)}+c_{w})^{T}w$$

still provides a statistically valid lower bound for $\phi(w)$.

We now explore the limiting behavior of $h_m(\cdot)$, which embeds all CP information accumulated over time.

Lemma 1. *The sequence $\{h_m(\cdot)\}_{m=1}^{\infty}$ converges uniformly on $\mathcal{W}$.*

Proof: $\mathbb{V}_m \subseteq \mathbb{V}_{m+1} \subseteq \mathbb{V}$ implies that $h_m(w,d_\xi) \leq h_{m+1}(w,d_\xi) \leq h(w,d_\xi)$ for all $w \in \mathcal{W}$ and d_ξ. Since $\{h_m(\cdot)\}_{k=1}^{\infty}$ increases monotonically and is bounded by the finite function $h(\cdot)$, it follows that $\{h_m(\cdot)\}_{k=1}^{\infty}$ converges point-wise to some function $\varphi(\cdot)$ satisfying $\varphi(w,d_\xi) \leq h(w,d_\xi)$ for all $w \in \mathcal{W}$, $d_\xi \in \Omega'$. Since $\mathbb{V}_m \subseteq \mathbb{V}_{m+1} \subseteq \mathbb{V}$, we have $\bar{\mathbb{V}} = \lim_{m\to\infty}\mathbb{V}_m \subseteq \mathbb{V}$. Since $\mathbb{V}$ is a finite set, so $\bar{\mathbb{V}}$ is also a finite set, and we have that:

$$\begin{aligned}\varphi(w,d_\xi) &= \lim_{m\to\infty} h_m(w,d_\xi) \\ &= \lim_{m\to\infty}\max\{\pi^T(r_\xi - Tw)|\pi\in\mathbb{V}_m\} \\ &= \max\{\pi^T(r_\xi - Tw)|\pi\in\bar{\mathbb{V}}\},\end{aligned}$$

and it follows that $\varphi(\cdot)$ is a continuous function. As $\mathcal{W}\times\Omega'$ is a compact space, and $\{h_m(\cdot)\}_{m=1}^{\infty}$ is a monotonic sequence of continuous functions, it follows that it converges uniformly to $\varphi(\cdot)$ (see Theorem 7.13 in Rudin (1976)). □

We now show that the sequence of CPs generated provides support points for $\phi(\cdot)$ in

the limit $m \to \infty$.

Theorem 5.3.3. *Let $\{w_{m_k}\}_{k=1}^{\infty}$ be an infinite subsequence of $\{w_m\}_{m=1}^{\infty}$. If $w_{m_k} \to \bar{w}$ then, with probability one,*

$$\lim_{k\to\infty} \frac{1}{m_k} \sum_{\xi=1}^{m_k} (\pi_{\xi}^{m_k})^T (r_{\xi} - T w_{m_k}) = \mathbb{E}[h(\bar{w}, D)].$$

In addition, every limit of $\{\alpha_{m_k}^{m_k}, \beta_{m_k}^{m_k} + c_w\}_{k=1}^{\infty}$ defines a support of $\phi(w)$ at $\bar{w}$, with probability one.

Proof: By definition, we have that

$$\begin{aligned} h_{m_k}(w_{m_k}, d_{\xi}) &= (\pi_{\xi}^{m_k})^T (r_{\xi} - T w_{m_k}) \\ \frac{1}{m_k} \sum_{\xi=1}^{m_k} h_{m_k}(w_{m_k}, d_{\xi}) &= \frac{1}{m_k} \sum_{\xi=1}^{m_k} (\pi_{\xi}^{m_k})^T (r_{\xi} - T w_{m_k}). \end{aligned}$$

By Lemma 1, there exists $\varphi(\cdot) \leq h(\cdot)$ such that $\{h_m\}_{m=0}^{\infty}$ converges uniformly to $\varphi(\cdot)$. We thus have:

$$\lim_{k\to\infty} \frac{1}{m_k} \sum_{\xi=1}^{m_k} \left[h_{m_k}(w_{m_k}, d_{\xi}) - \varphi(\bar{w}, d_{\xi}) \right] = 0,$$

and $\lim_{k\to\infty} \frac{1}{m_k} \sum_{\xi=1}^{m_k} h(w, d_{\xi}) = \mathbb{E}[h(w, D)]$ with probability one. It is now sufficient to show that $\varphi(\bar{w}, d_{\xi}) = h(\bar{w}, d_{\xi})$ with probability one. Let d_{ξ} be a given realization and suppose that, for every $\delta > 0$, we have $\mathbb{P}\{|D - d_{\xi}| < \delta\}$. Then, for every $\delta > 0$, $|d_{m_k} - d_{\xi}| < \delta$ infinitely often, with probability one. Because $h(\cdot)$ is a continuous function and $\{h_m(\cdot)\}_{m=1}^{\infty}$ is uniformly convergent, for every $\epsilon > 0$ there exists a $\delta > 0$ and $N < \infty$ such that $|(\bar{w}, d_{\xi}) - (w, d)| < \delta$ and with:

$$\begin{aligned} &|h(\bar{w}, d_{\xi}) - h(\bar{w}, d)| < \frac{\epsilon}{3} \\ &|h_{m_k}(\bar{w}, d_{\xi}) - h_{m_k}(\bar{w}, d)| < \frac{\epsilon}{3}, \quad k \geq N. \end{aligned}$$

Consequently, because $\lim_{k\to\infty} w_{m_k} = \bar{w}$, we have $\mathbb{P}\{|D - d_{\xi}| < \delta\}$ implies that for every

$\epsilon > 0$ there exists a further subsequence $\{(w_{m'_k}, d_{m'_k})\}_{k=1}^{\infty}$ and $K < \infty$ such that

$$|h(\bar{w}, d_\xi) - h_{m'_k}(\bar{w}, d_{m'_k})| < \frac{\epsilon}{3}$$
$$|h(\bar{w}, d_{m'_k}) - h(w_{m'_k}, d_{m'_k})| < \frac{\epsilon}{3}$$
$$|h_{m'_k}(w_{m'_k}, d_{m'_k}) - h_{m'_k}(w_{m'_k}, d_\xi)| < \frac{\epsilon}{3}$$

for all $m'_k \geq K$. By construction, $h_{m'_k}(w_{m'_k}, d_{m'_k}) = h(w_{m'_k}, d_{m'_k})$. Thus, for every $\epsilon > 0$, there exists a subsequence $\{w_{m'_k}\}_{k=1}^{\infty}$ and $K < \infty$ such that

$$|h(\bar{w}, d_\xi) - h_{m'_k}(w_{m'_k}, d_\xi)| < \epsilon$$

for all $m'_k \geq K$. Consequently, it follows that $\varphi(\bar{w}, d_\xi) = h(\bar{w}, d_\xi)$. Finally, since Ω' is compact, we have that $\mathbb{P}\{|D - d_\xi| < \delta\}$ for some $\delta > 0$ and for only finitely many values of d_ξ, with probability one. Thus, with probability one, $\varphi(\bar{w}, d_\xi) = h(\bar{w}, d_\xi)$ for all but a finite number of realizations. Hence, with probability one,

$$\lim_{k\to\infty} \frac{1}{m_k} \sum_{\xi=1}^{m_k} (\pi_\xi^{m_k})^T (r_\xi - T w_{m_k}) = \mathbb{E}[h(w, D)].$$

Moreover, since $h(w, d_\xi) = \max\{\pi^T(r_\xi - Tw) | \pi \in \mathbb{V}\}$ and $\mathbb{V}_m \subset \mathbb{V}$ for all m, it follows that

$$\begin{aligned} c_w^T w + \frac{1}{m_k} \sum_{\xi=1}^{m_k} h(w, d_\xi) &\geq c_w^T w + \frac{1}{m_k} \sum_{\xi=1}^{m_k} \pi_\xi^{m_k T} (r_\xi - Tw) \\ &= \alpha_{m_k}^{m_k} + \left(\beta_{m_k}^{m_k} + c_w\right)^T w. \end{aligned}$$

We conclude that, with probability one, $\phi(w)$ is at least as large as any limiting cut that is associated with the subsequence of cuts defined by $\{(\alpha_{m_k}^{m_k}, \beta_{m_k}^{m_k} + c)\}_{k=1}^{\infty}$. Thus, any limiting cut defines a support of $\phi(w)$ at $\bar{w}$. $\square$

Theorem 5.3.4. *There exists a subsequence of* $\{w_{m_k}\}_{k=1}^{\infty}$, *satisfying* $\lim_{k\to\infty}(\underline{\phi}_{m_k}(w_{m_k}) -$

$\underline{\phi}_{m_k-1}(w_{m_k})) = 0.$

Proof: See proof of Theorem 3 in Higle and Sen (1991).

We now establish the main convergence result.

Theorem 5.3.5. *There exists a subsequence $\{w_{m_k}\}_{k=1}^{\infty}$, such that every accumulation point of $\{w_{m_k}\}_{k=1}^{\infty}$ is an optimal solution w^* of $\mathbf{P}_\infty$, with probability one.*

Proof: Let w^* be an optimal solution of $\mathbf{P}_\infty$. From Theorem 5.3.4, there exists a subsequence $\{w_{m_k}\}_{k=1}^{\infty}$ such that $\lim_{k\to\infty}(\underline{\phi}_{m_k}(w_{m_k}) - \underline{\phi}_{m_k-1}(w_{m_k})) = 0$. Let $\{w_{m_k}\}_{k\in\mathcal{K}}$ be a further subsequence such that $\lim_{k\in\mathcal{K}} w_{m_k} = \bar{w}$. Compactness of $\mathcal{W}$ ensures that $\bar{w} \in \mathcal{W}$, and thus $\phi(w^*) \leq \phi(\bar{w})$. We know that:

$$\underline{\phi}_m(w) \leq c_w^T w + \frac{1}{m}\sum_{\xi=1}^{m} h(w, d_\xi),\ w \in \mathcal{W}, \tag{5.3.11}$$

and thus,

$$\limsup_{m\in\mathcal{M}} \underline{\phi}_m(w^*) \leq c_w^T w^* + \mathbb{E}[h(w^*, D)] = \phi(w^*) \tag{5.3.12}$$

with probability one for any index set $\mathcal{M}$. Since w_m minimizes $\underline{\phi}_{m-1}(\cdot)$ on $\mathcal{W}$, we have $\underline{\phi}_{m-1}(w_m) \leq \underline{\phi}_{m-1}(w^*)$ for all m. From Theorem 5.3.3, $\lim_{k\in\mathcal{K}} \phi_{m_k}(w_{m_k}) \leq \phi(\bar{w})$ with probability one and, also by definition, $\lim_{k\in\mathcal{K}} \underline{\phi}_{m_k}(w_{m_k}) - \underline{\phi}_{m_k-1}(w_{m_k}) = 0$. Thus, we have $\lim_{k\in\mathcal{K}} \underline{\phi}_{m_k-1}(w_{m_k}) = \phi(\bar{w})$, with probability one. Combining these results we obtain:

$$\begin{aligned}
\phi(w^*) &\leq \phi(\bar{w}) \\
&= \limsup_{k\in\mathcal{K}} \underline{\phi}_{m_k-1}(w_{m_k}) \\
&\leq \underline{\phi}_{m_k-1}(w^*) \leq \phi(w^*),
\end{aligned}$$

with probability one. We thus have that $\phi(\bar{w}) = \phi(w^*)$. □

5.3.3 *Short-Term MPC Controller*

The CP scheme is guaranteed to deliver optimal targets as data is accumulated over time. Notably, because the scheme is inherently retroactive, it achieves optimal targets without using any data forecasts. So the question is: How does the proposed scheme accommodate forecast information? From an implementation stand-point, another important question is: What metrics can one use to monitor the optimality of the targets?

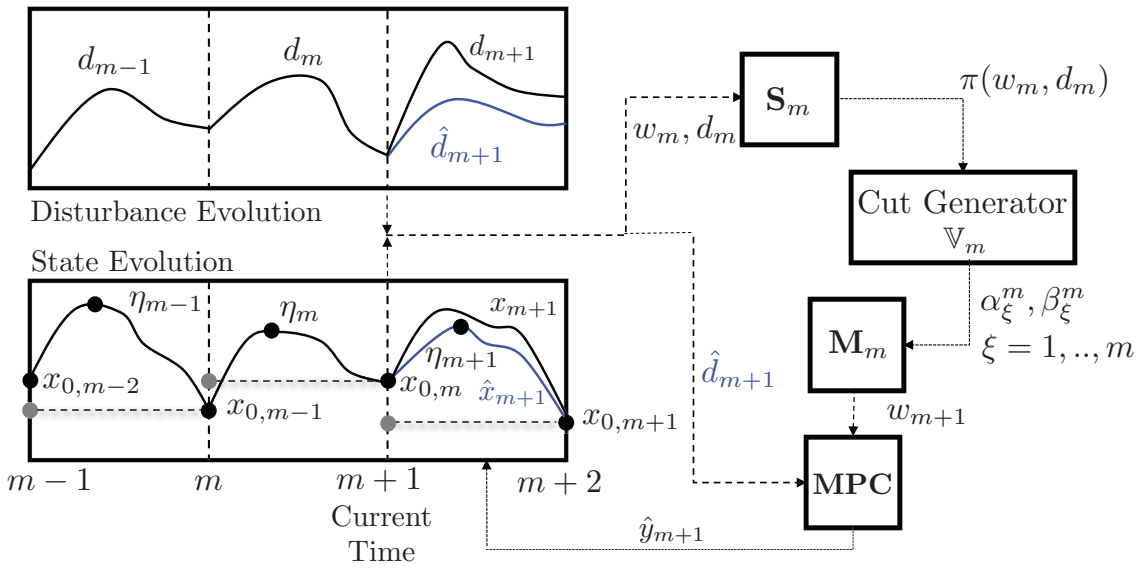

Figure 5.1: Sketch of hierarchical scheme using cutting planes.

The proposed scheme offers a couple of mechanisms to embed forecast information. First, initial guesses for the targets $w_m = (x_{m,0}, \eta_m)$ at period $m = 1$, can be obtained by solving $\mathbf{P}_{m'}$ for some $m' \geq 1$ that uses a data forecast $\{\hat{d}_\xi\}_{\xi=1}^{m'}$. Because the forecast is expected to contain errors, the initial targets are expected to be suboptimal but these will be refined as true data is obtained. Also, as discussed in Section 5.3.1, the proposed scheme also enables the incorporation of forecasts at the beginning of each period to compute the intra-period policies. In particular, given that the system is at x_m, the current guess for the targets w_{m+1}, and a forecast $\hat{d}_{m+1}$ over period $m+1$, one can compute the internal control and state policies $\hat{y}_{m+1} = (\hat{x}_{m+1}, \hat{u}_{m+1})$ that satisfy the targets w_{m+1} using an intra-period

MPC controller. In the linear case, this can be done by solving:

$$\min_{y \in \mathbb{R}_+^{n_y}} \hat{c}_{m+1}^T y \text{ s.t. } Wy = \hat{r}_{m+1} - Tw_{m+1}. \tag{5.3.13}$$

Clearly, the policies $\hat{y}_{m+1}$ are suboptimal because the forecast $\hat{d}_{m+1}$ will deviate from the true realization d_{m+1} (this becomes known at the end of period $\xi + 1$). Because the cost function $h(\cdot, \cdot)$ is continuous, one can use standard perturbation results to show that the optimality error in the intra-period policy is bounded by the forecast error as $|h(w_{m+1}, d_{m+1}) - h(w_{m+1}, \hat{d}_{m+1})| \leq L\|d_{m+1} - \hat{d}_{m+1}\|$ for some Lipschitz constant $L \in \mathbb{R}_+$ (see Theorem 4.156 in Bonnans and Shapiro (2013)). This implies that the quality of the forecast does affect optimality with respect to the intra-period policies. Interestingly, however, the short-term MPC controller only needs to have a forecast over a period of length n (as opposed to over the entire horizon p). Consequently, the hierarchical MPC controller is less sensitive to forecast errors than standard proactive MPC approaches. In addition, we emphasize that the forecast quality does not affect optimality of the targets.

The period length n introduces an interesting and important trade-off between long-term and short-term performance. As discussed previously, increasing n ensures that the resulting periodic policy better approximates the policy of problem $\mathbf{O}_m$. On the other hand, increasing n indicates that the short-term MPC controller needs to run over a longer period; as a result, it is more computationally expensive and susceptible to forecast errors. For instance, as we have noted, as n is increased the hierarchical MPC scheme resembles a standard proactive MPC scheme (since the periodicity constraints are enforced less often).

The full hierarchical scheme is sketched in Figure 5.1 and is summarized as follows:

1. Initialize $m \leftarrow 1$, $\mathbb{V}_m \leftarrow \emptyset$, and w_m.

2. Use forecast $\hat{d}_m$ and targets w_m to obtain intra-period policies $\hat{y}_m$ using MPC controller.

3. Let system transition from $m \rightarrow m + 1$.

4. Observe d_m and solve $\mathbf{S_m}$ to obtain $\pi(w_m, d_m)$.

5. Update $\mathbb{V}_m \leftarrow \mathbb{V}_{m-1} \cup \{\pi(w_m, d_m)\}$.

6. Obtain $\alpha_\xi^m, \beta_\xi^m$, $\xi = 1, \ldots, m$ from cut generator.

7. Solve $\mathbf{M}_m$ and obtain updated targets w_{m+1}.

8. Shift period time $m \leftarrow m + 1$ and return to Step 2.

The SP interpretation allows us to derive metrics and techniques to monitor optimality. We first note that the running cost $\phi_m(w_m)$ evaluated at the current target w_m can be evaluated by solving the sequence of subproblems $\mathbf{S}_\xi$, $\xi = 1, \ldots, m$. The running cost is an upper bound of the optimal running cost (obtained by solving $\mathbf{P}_m$). Moreover, the proposed CP scheme offers the guarantee that the cost $\underline{\phi}_m(w_m)$ is a lower bound of the running cost $\phi_m(w_m)$, which is an asymptotically exact statistical approximation of $\phi(w_m)$. The difference between the running cost (upper bound) and the lower bound is known in the SP literature as the optimality gap and is formally defined as $\epsilon_m := (\phi_m(w_m) - \underline{\phi}_m(w_m))/\phi_m(w_m)$. Here, we refer to ϵ_m as the current gap. The convergence of the CP scheme guarantees that $\lim_{m\to\infty} \epsilon_m = 0$. We note that this gap can be used to measure the quality of the CP approximation but should be used with care when interpreting optimality. In particular, the gap can only be used as a measure of optimality in the limit $m \to \infty$ (once the running cost $\phi_m(w_m)$ is close to the actual cost $\phi_\infty(w_m)$). Consequently, one usually resorts to computing confidence intervals for $\phi_m(w_m)$ by using inference (evaluate the cost at w_m but using different combinations of realizations) (Linderoth et al., 2006). Motivated by this, in benchmarking studies, we are also interested in monitoring the overall optimality gap $\bar{\epsilon}_m := (\phi_\infty(w_m) - \underline{\phi}_m(w_m))/\phi_\infty(w_m)$, provided that $\phi_\infty(w_m)$ can be computed.

5.3.4 *Extensions to Nonlinear Systems*

In the case of linear systems, the CP scheme can approximate the running cost $\phi_m(w)$ using a finite number of supporting hyperplanes, which keeps the master problem $\mathbf{M}_m$ tractable.

The SP representation opens the door to other schemes such as proximal point methods for nonlinear (but convex) problems (Bertsekas, 2011). Here, the idea is to prevent the accumulation of data over time by summarizing past information in terms of a proximal term. In our context, for instance, the proximal point strategy will result in a problem of the form at stage m:

$$\min_{w \in \mathcal{W}} \ \mu \|w - w_m\|^2 + g(w) + \frac{1}{m'} \sum_{\xi=m'}^{m} h(w, d_\xi). \tag{5.3.14}$$

Here, $\mu \|w - w_m\|^2$ is a regularization term with parameter $\mu > 0$. This term summarizes data before time m' and resembles the arrival cost used in moving horizon estimation (Rao et al., 2003). In the general case of nonconvex problems, one can use a statistical approximation $\min_{w \in \mathcal{W}} \ \phi_m(w)$ at every period m by using linear algebra decomposition schemes. In particular, it is well-known that problems with the structure of $\mathbf{P}_m$ give linear algebra systems that enable parallel decomposition (Zavala et al., 2008). Given that the coupling is only in the space of the periodic targets x_0, this approach can scale to systems with thousands of states and tens of thousands of periods (realizations). This approach, however, exhibits a fundamental limitation in the number of periods that it can handle (because data is accumulated unboundedly over time). This is, in fact, also a limitation of statistical approximation schemes for SP. To circumvent this issue, one can use clustering techniques that seek to compress the realization space to maintain a tractable approximation (Cao et al., 2016). Such techniques are based on the observation that data realizations tend to be redundant and only a small subset actually impacts the cost. One can quantify the error incurred in the scenario compression by using inference techniques (Linderoth et al., 2006).

5.3.5 *Stability Considerations*

The proposed hierarchical schemes provide important stability properties. A formal treatment of such properties is beyond the scope of this work, but here we present some basic arguments. We first consider the *nominal case*, in which the data realization in each period

is the same ($d_{\xi} = d_{\xi'}$ for all ξ, ξ') and it is known. This is equivalent to assuming that the dynamics are of the form $x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t})$ (they are time-invariant). We also assume that no peak cost $\phi_2(\cdot)$ is used. This nominal setting is considered in the periodic MPC formulations of Huang et al. (2011) andZanon et al. (2017). Both of these formulations use a proactive RH strategy and enforce a periodicity constraint at the end of the horizon. In Huang et al. (2011), it is shown that their MPC scheme delivers a periodic state that is a steady-state (the closed-loop system is stable). In Zanon et al. (2017), it is shown that their MPC scheme converges to a periodic state (not necessarily a steady-state), and is thus stable in this sense. Converse to a periodic state that is not a steady-state is desirable because this provides flexibility. Both of these proactive schemes require dissipativity properties. For this same nominal setting, we note that the solution of P_∞ also delivers a periodic state (by construction). Since the proposed hierarchical scheme converges to the periodic state of P_∞, we thus have that our approach is stable in this sense. This result does not require any dissipativity assumptions.

For the more general case in which the data changes in each period, we have that the dynamics are time-variant. Moreover, we have that the data cannot be forecast perfectly. Surprisingly, for this more challenging setting, we have that the retroactive scheme proposed also delivers the optimal periodic state and thus is stable. This is a remarkable result that no other known scheme reported in the literature provides.

The short-term MPC controller admits a standard state-feedback control representation. This is because the control policy inside the period is updated based on the current state and the desired state target. The long-term controller (updating the targets), however, does not admit such a representation because it is retroactive (and not proactive). In particular, it uses the entire history (and not just the current state) to compute the next targets. In future work, we will formalize the stability analysis of the approach and the associated state-feedback representation.

5.4 Computational Experiments

The performance of the proposed scheme is demonstrated using an application in buildings with electricity storage. The goal of the controller is to determine optimal short-term (hourly) participation strategies in frequency regulation (FR) and energy markets that maximize revenue and simultaneously mitigate long-term demand charges. We consider a utility-scale stationary battery with a capacity of 0.5 MWh, rated power of 1 MW, and a ramping limit of 0.5 MW/hr. We use real data for energy and FR prices from PJM Interconnection (shown in Figure 5.2). We also use real load data for a typical university campus (shown also in Figure 5.2). The periodic components in the energy prices and the load profile can be clearly observed, while the periodicity of FR prices is not as strong.

The MPC problem is formulated using daily periods of $n = 24$ hours and we consider an horizon of $m = 300$ days (nearly a year). The model parameters include the battery storage capacity ($\overline{E}$ kWh), maximum discharging and charging rates (power) ($\overline{P}, \underline{P}$ in kW, respectively), minimum fraction of battery capacity reserved for frequency regulation (ρ in kWh/kW), and maximum ramping limit ($\overline{\Delta P}$ in kW/h). The random data are the loads ($L_{\xi,t}$ kW), market prices for electricity and FR capacity ($\pi^e_{\xi,t}$ \$/kWh and $\pi^f_{\xi,t}$ \$/kW respectively), demand charge (monthly) (π^D in \$/kW), and fraction of FR capacity requested by ISO ($\alpha_{\xi,t}$). The decision variables are net battery discharge rate (power) ($P_{\xi,t}$ in kW), FR capacity provided to ISO ($F_{\xi,t}$ in kW), state-of-charge (SOC) of the battery ($E_{\xi,t}$ in kWh), load requested from utility ($d_{\xi,t}$ in kW) and peak load (D in kW). The formulation minimizes the total cost (negative total revenue), which is given by the demand charge (peak cost) and the revenues collected from the market (time-additive cost). Detailed notation and analysis of this formulation is presented in Appendix B and in Kumar et al. (2018d).

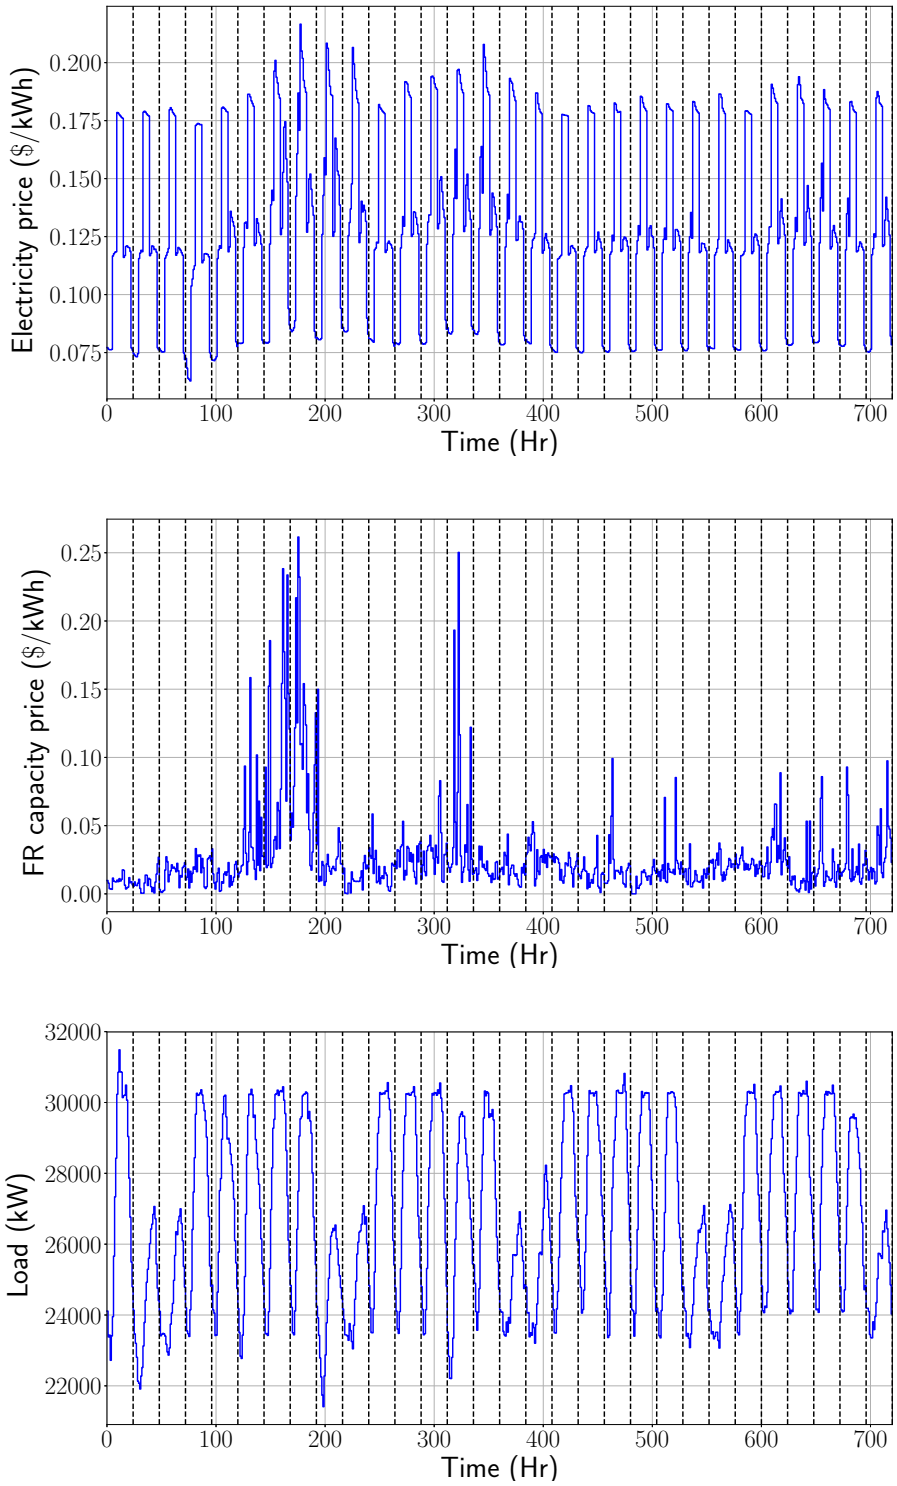

Figure 5.2: Energy price (top), FR price (middle), and load (bottom) data.

The problem has the form:

$$\min \frac{1}{m} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}} \left(-\pi^e_{\xi,t}(P_{\xi,t} - \alpha_{\xi,t} F_{\xi,t}) - \pi^f_{\xi,t} F_{\xi,t} \right) + \pi^D D$$

$$\text{s.t. } P_{\xi,t} + F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}, \xi \in \Xi$$

$$P_{\xi,t} - F_{\xi,t} \geq -\underline{P},\ t \in \mathcal{T}, \xi \in \Xi$$

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \bar{\mathcal{T}}, \xi \in \Xi$$

$$\rho F_{\xi,t} \leq E_{\xi,t} \leq \overline{E} - \rho F_{\xi,t},\ t \in \mathcal{T}, \xi \in \Xi$$

$$\rho F_{\xi,t} \leq E_{\xi,t+1} \leq \overline{E} - \rho F_{\xi,t},\ t \in \bar{\mathcal{T}}, \xi \in \Xi$$

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}, \xi \in \Xi$$

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \mathcal{T}, \xi \in \Xi$$

$$d_{\xi,t} \leq D,\ t \in \mathcal{T}, \xi \in \Xi$$

$$P_{\xi,t} + F_{\xi,t} \leq L_{\xi,t},\ t \in \mathcal{T}, \xi \in \Xi$$

$$E_{\xi,0} = E_0,\ \xi \in \Xi$$

$$E_{\xi,N_\xi} = E_0,\ \xi \in \Xi$$

$$0 \leq E_{\xi,t} \leq \overline{E},\ t \in \mathcal{T}, \xi \in \Xi$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}, \xi \in \Xi$$

$$0 \leq F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}, \xi \in \Xi$$

We first solve the long-horizon MPC problem (5.2.1) for a horizon of 300 days assuming perfect knowledge of the data to obtain the optimal policy. We compare this policy against that of a long-horizon MPC formulation that enforces periodicity constraints (5.2.2). The optimal and periodic policies over 30 days are presented in Figure 5.3. Here, we can see that the policies match. The total cost of the the long-horizon MPC problem is \$136,050 while the cost of the long-horizon problem with periodicity constraints is \$136,068. We can see that, in this application, state periodicity arises naturally because the battery needs to maintain a minimum SOC level after each period. We then ran the proposed retroac-

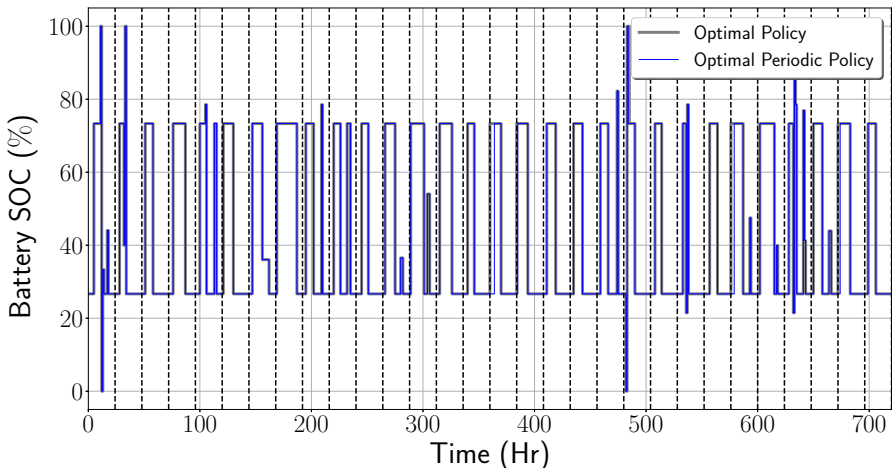

Figure 5.3: Optimal and optimal periodic policies.

tive CP scheme for 300 periods to progressively update the targets. The evolution of the current gap ϵ_m and overall optimality gap $\bar{\epsilon}_m$ is shown in Figure 5.4. We observe that the overall gap eventually vanishes, demonstrating that the CP scheme delivers optimal targets. The overall gap closes to zero close to the end of the horizon, once the peak demand is observed. We also see that the current gap closes to 0% in about 50 periods and stays there for the rest of the horizon. This illustrates that the cutting planes approximate the

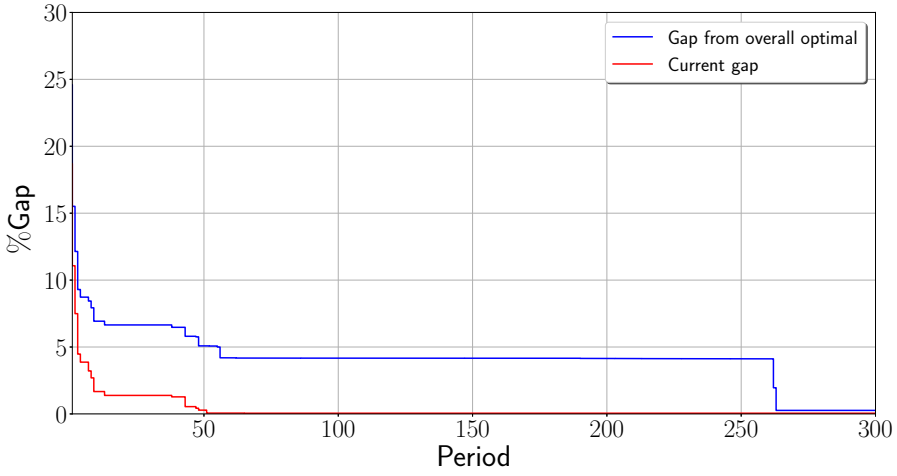

Figure 5.4: Evolution of optimality gap.

cost function well, but also that this metric can be misleading. We have also found that the performance of the proposed hierarchical scheme is close to that obtained with the optimal long-term periodic policy. In particular, the cost of the hierarchical MPC is \$139,978 (a difference of 2.89%). Accelerating the convergence of the CP scheme can help decrease this gap. We also compared the performance of the hierarchical scheme with that of a standard (non-periodic) MPC scheme that uses a prediction horizon one and fourteen days. The corresponding costs are \$139,884 and \$138,739. We can see that the cost decreases slowly with

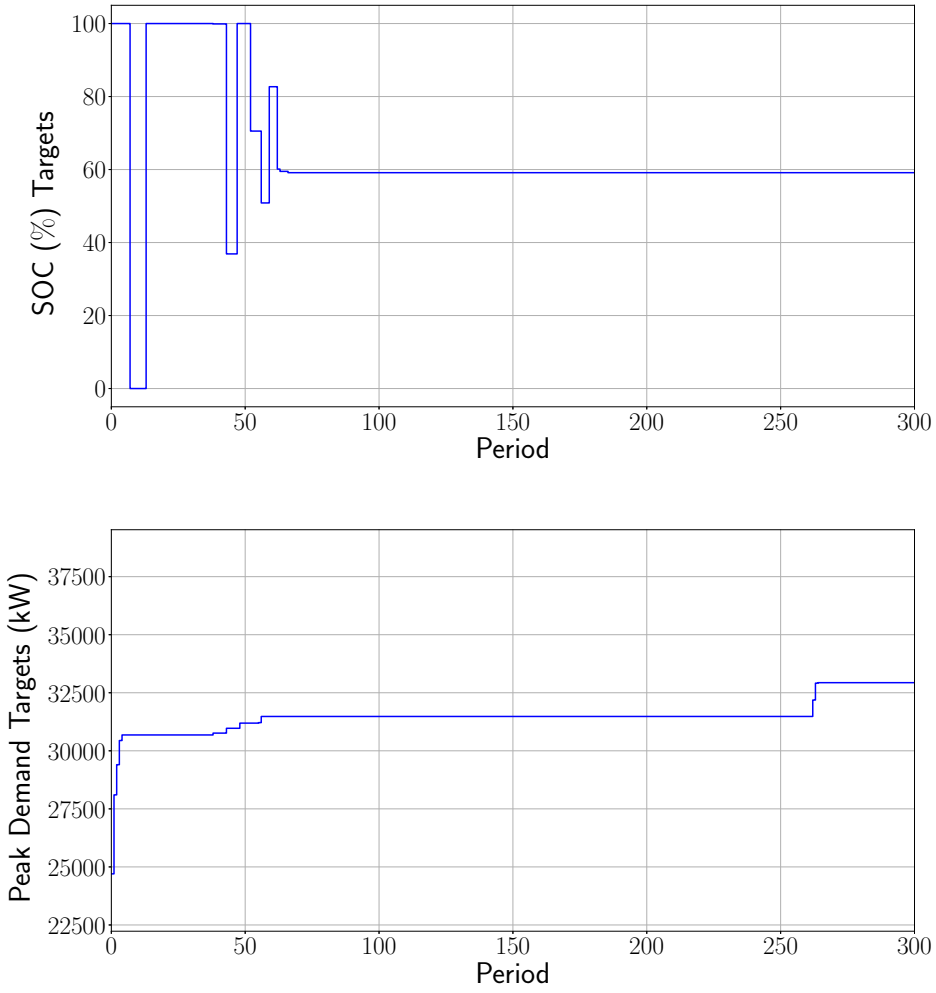

Figure 5.5: Evolution of periodic SOC (top) and peak (bottom) targets obtained with cutting-plane scheme.

the prediction horizon. We also observe that the cost for standard MPC with a one-day prediction horizon is similar to that of the hierarchical scheme. The hierarchical scheme, however, offers optimality guarantees (standard MPC does not).

Figure 5.5 shows the evolution of the periodic state and peak targets. We see that the CP scheme adaptively updates the targets as data is accumulated over time. The SOC target settles quickly to the optimal level of 59%. The peak target requires more periods to settle and this behavior is attributed to the fact that the peak load is observed at period 264. After this period the peak settles at its optimal value of 32,935 kW. Figure 5.6 shows the intra-period policies for the short-term MPC controller for the first 7 days (periods) of operation. We see that the controller follows the target of the CP scheme.

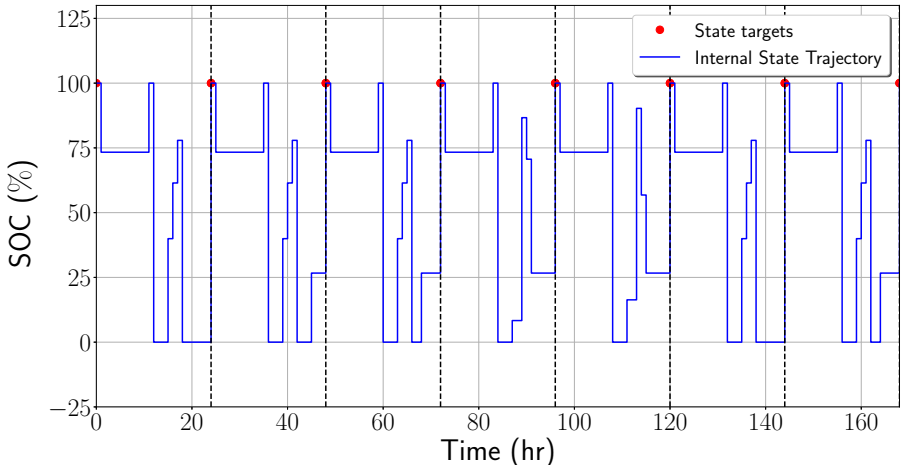

Figure 5.6: Evolution of periodic SOC targets and intra-period policies obtained with CP scheme (first seven periods).

We also compared the retroactive hierarchical MPC scheme with the proactive MPC approach for periodic systems of Huang et al. (2011). For this scheme, we consider a prediction horizon of 7 days and a period of 1 day. For this comparison, we removed the demand charge (peak cost) from the formulation and assumed that disturbances can be forecast perfectly. In Figure 5.7, we present the evolution of the periodic SOC for both approaches. We observe that, with the proactive approach, the periodic state does not

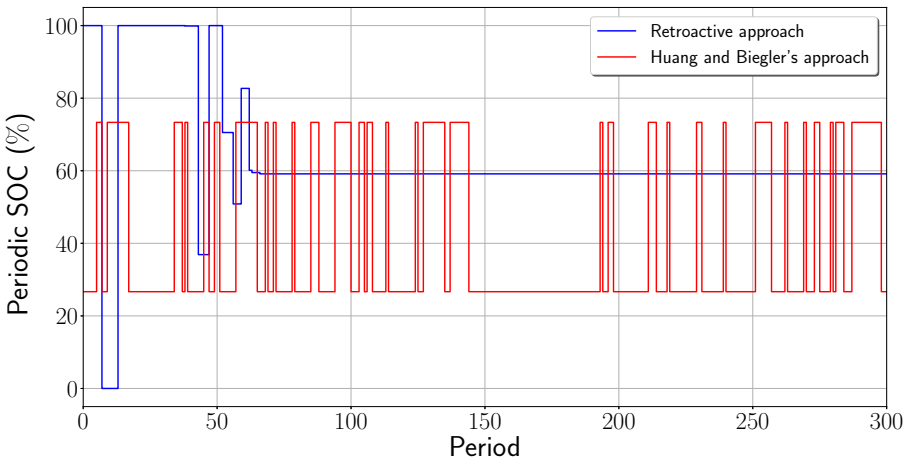

Figure 5.7: Comparison of evolution of periodic SOC resulting from the retroactive approach and proactive MPC approach of Huang et al. (2011).

converge. This is because the dynamics are time-varying. A similar behavior would be expected from the proactive scheme of Zanon et al. (2017). On the other hand, the retroactive approach converges to a periodic state. Moreover, we found that the cost of the proactive scheme is 5% worse than that of the retroactive scheme. This is because the proactive scheme does not capture the long-term trend of the disturbances. In other words, the retroactive scheme delivers a policy that is close-to-optimal.

5.5 Conclusions

In this chapter, we showed that stochastic programming provides a framework to design hierarchical MPC schemes for periodic systems. We have shown that, under periodicity, it is possible to compute and refine periodic state targets by solving a retroactive optimization problem that progressively accumulates historical data. The retroactive problem is a statistical approximation of the stochastic program that delivers optimal targets in the long run to guide a short-term MPC controller. The computation of the optimal targets can be achieved without data forecasts. The SP setting opens the door to a number of potential developments in hierarchical MPC scheme.

6

A STOCHASTIC DUAL DYNAMIC PROGRAMMING FRAMEWORK FOR MULTISCALE MPC

In this chapter, we provide a scalable approach to handle complex MPC applications with uncertainties evolving over long time horizons and with fine time resolutions by deriving and interpreting stochastic dual dynamic programming (SDDP) from the perspective of MPC. Scalability is enabled by the use of a nested cutting plane scheme, which uses forward and backward sweeps along the time horizon to adaptively construct and refine cost-to-go functions (Kumar et al., 2018c). SDDP can also select random paths on-the-fly, thus avoiding the exponential complexity associated with the construction of scenario trees and enabling the construction of statistical performance bounds. Our work seeks to establish connections between SDDP and MPC that allow us to obtain performance bounds for deterministic and stochastic MPC and to propose new MPC schemes that can deal with multiple timescales. We demonstrate the developments using a stationary battery system that participates in multiscale energy markets.

6.1 Introduction

Scalability bottlenecks of MPC due to long time horizons and fine time resolutions are well known. These issues are drastically exacerbated in stochastic MPC formulations, which need to deal with uncertainty characterizations that capture long horizons and fine resolutions. Specifically, uncertain disturbances and associated state/control policies often need to be represented in the from scenario trees, which grow exponentially in size with the number of time steps (Birge and Louveaux, 2011). Receding horizon implementations of two-stage stochastic MPC formulations are often used to compute approximate policies (Kumar et al., 2018a). Scalability can also be enabled by using robust formulations, which embed specialized policies. It is well-known, however, that such formulations can only handle certain types of uncertainty descriptions (Diehl and Bjornberg, 2004).

Several strategies have been devised in the stochastic programming community to deal with large-scale multi-stage stochastic programs (SP) (Shapiro and Wardi, 1996). Of particular interest is a technique called stochastic dual dynamic programming (SDDP), which is a decomposition approach originally devised to handle linear multi-stage SPs arising in energy planning (Pereira, 1989). SDDP is a nested Benders decomposition scheme which progressively approximates the recourse cost function by using cutting planes. When analyzed from a dynamic programming perspective, one can show that the recourse function is the so-called cost-to-go. In SDDP, state policies and cutting plane information are collected by using forward and backward sweeps to refine the cost-to-go functions. Notably, SDDP enables *on-the-fly sampling* of scenario tree paths and thus avoids the need to form the multi-stage SP explicitly (which is often impossible due to the exponential growth in memory requirements). When the realizations are independent, one can use a compact scenario tree in which the number of nodes scales linearly with the number of stages. This compact tree representation encodes information on all possible paths for the random disturbances. Figure 6.1 shows the actual scenario tree in which the number of nodes grows exponentially with the number of stages and a compact representation of the scenario tree.

This is a key benefit over other decomposition methods for multistage SP such as interior-point and Lagrangian dual decomposition (Kouzoupis et al., 2018), which need to explicitly form and store the extensive form of the SP. SDDP is thus the only known paradigm that can handle SPs with many stages and thus is suitable to handle problems with long time horizons and fine time resolutions. On the other hand, finite convergence of SDDP can only be established for linear systems (Shapiro, 2011).

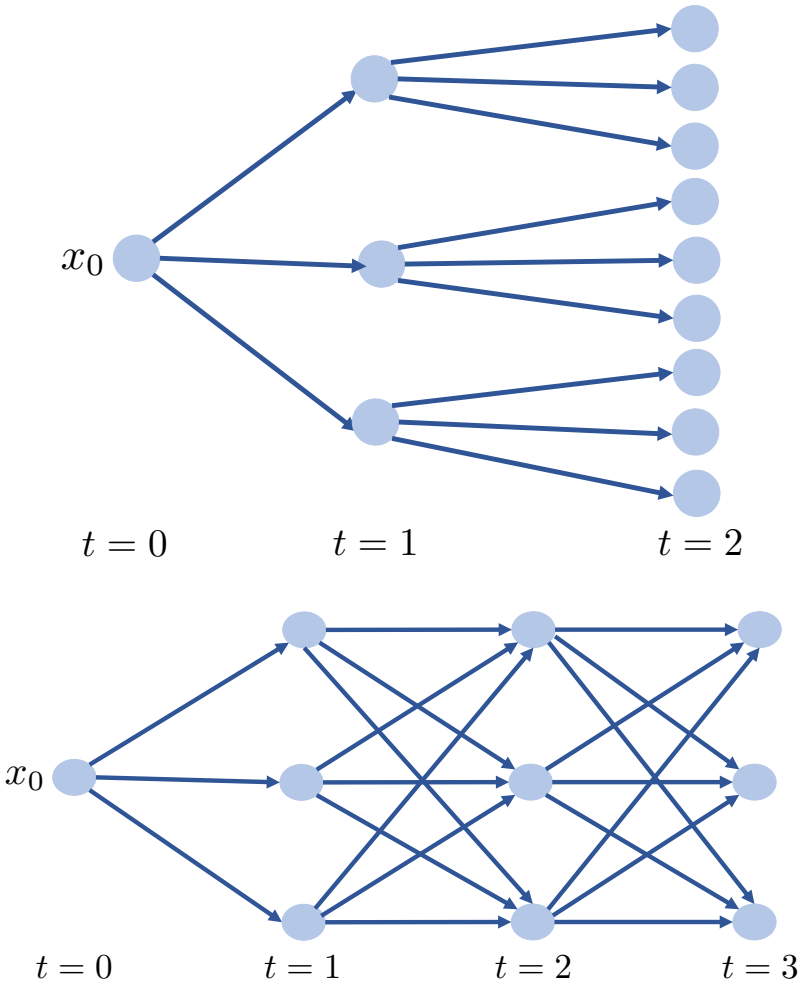

Figure 6.1: Full (top) and compact (bottom) scenario tree representations for multi-stage SP.

In this work, we present a formal derivation and interpretation of SDDP from the perspective of MPC. In particular, we derive SDDP for state-space representations with ad-

ditive disturbances in a deterministic and a stochastic setting. This establishes interesting and important connections that allow us to obtain bounds in performance for deterministic and stochastic MPC and to propose MPC schemes that can deal with multiple timescales. In particular, we show that a forward SDDP sweep can be interpreted as a receding horizon MPC scheme applied over a long horizon and that a backward sweep collects adjoint information to construct an approximation of the cost-to-go. This provides a mechanism to construct terminal costs for short-term MPC controllers and thus enables the construction of hierarchical MPC architectures. A similar observation was recently made in the context of block Gauss-Seidel (GS) decomposition schemes applied to MPC (which can also be interpreted as a forward-backward scheme applied to the Euler-Lagrange conditions) (Zavala, 2016; Shin and Zavala, 2018). As in the GS approach, SDDP solves a stage subproblem to minimize the current stage cost and a terminal cost that uses adjoint information (that captures the price of future information). Compared to GS, however, one can show that SDDP is convergent and provides lower bounds of the optimal objective value (because the cost-to-go is constructed from supporting hyperplanes). Moreover, one can show that the running average of the cost (over all iterations) provides an upper bound. These bounding properties can be used to establish cost bounds for receding horizon MPC policies. We propose a modification to the standard SDDP scheme to handle more general MPC formulations that include multiple time steps per stage and uncertainties that are resolved at different timescales. Specifically, we outline how to properly construct cutting planes in such formulations to ensure convergence.

The proposed developments are motivated by energy systems and inventory management. In the context of energy systems, long planning horizons are often needed to capture monthly peak electricity costs (demand charges) (Ma et al., 2012a; Oudalov et al., 2007a; Kumar et al., 2018a), and to capture variations in demand, supply, and prices at various timescales (Ma et al., 2012a; Oudalov et al., 2007a; Braun, 1990).

6.2 Deterministic Multistage Formulation

We begin by considering the deterministic MPC problem:

$$\min_{\mathbf{x}_{\mathcal{N}}, \mathbf{u}_{\mathcal{N}}} \sum_{t=0}^{N-1} a_t^T x_t + b_t^T u_t \tag{6.2.1a}$$

$$\text{s.t. } x_{t+1} = Ax_t + Bu_t + Cd_t,\ t = 0, \ldots, N-1 \tag{6.2.1b}$$

$$u_t \in \mathcal{U},\ t = 0, \ldots, N-1 \tag{6.2.1c}$$

$$x_t \in \mathcal{X},\ t = 0, \ldots, N \tag{6.2.1d}$$

Note that we specifically consider a linear MPC problem in this work. Here, $a_t \in \mathbb{R}^{n_x}$ and $b_t \in \mathbb{R}^{n_u}$ are stage costs associated with the states $x_t \in \mathbb{R}^{n_x}$ and controls $u_t \in \mathbb{R}^{n_u}$, respectively. The system disturbances are denoted as $d_t \in \mathbb{R}^{n_d}$ and the horizon length is denoted as N. We denote trajectories for controls and states over the horizon $\mathcal{N} := \{0, \ldots, N\}$ as $\mathbf{u}_{\mathcal{N}} := (u_t, u_{t+1}, \ldots, u_{t+N-1})$ and $\mathbf{x}_{\mathcal{N}} := (x_{t+1}, x_{t+2}, \ldots, x_{t+N})$. The dynamics of the linear system are described by the matrices $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$, and $C \in \mathbb{R}^{n_x \times n_d}$. The state x_t is provided as initial condition. The states x_t and controls u_t are bounded by the polyhedral sets $\mathcal{X}$ and $\mathcal{U}$, respectively.

We assume that the disturbances are described by a stochastic process $D_{\mathcal{N}} := (D_0, D_1, \ldots, D_{N-1})$. We define the information (realizations) collected from process $D_{\mathcal{N}}$ up to time t as $I_t := (d_0, d_1, \ldots, d_t)$. We also denote a trajectory of realizations as $\mathbf{d}_{\mathcal{N}} := (d_0, d_1, \ldots, d_{N-1})$. We thus see that the deterministic problem (6.2.1) uses a realization of the process (obtained from a forecast).

We now consider a reformulation of a deterministic MPC problem as a dynamic program (in multi-stage form). In this formulation, we use a realization of the process $\mathbf{d}_{\mathcal{N}}$ and the objective function includes the stage cost for the current time and the cost-to-go

function which summarizes the cost of future stages. The problem takes the form:

$$Q_0(x_0, d_0) = \min_{\substack{x_1 \in \mathcal{X}, \\ u_0 \in \mathcal{U}}} \; a_0^T x_0 + b_0^T u_0 + Q_1(x_1, d_1) \tag{6.2.2a}$$

$$\text{s.t. } x_1 = Ax_0 + Bu_0 + Cd_0 \tag{6.2.2b}$$

where the cost-to-go is:

$$Q_1(x_1, d_1) = \min_{\substack{x_2 \in \mathcal{X}, \\ u_1 \in \mathcal{U}}} \; a_1^T x_1 + b_1^T u_1 + Q_2(x_2, d_2) \tag{6.2.3a}$$

$$\text{s.t. } x_2 = Ax_1 + Bu_1 + Cd_1. \tag{6.2.3b}$$

Following the recursion we obtain the cost-to-go:

$$Q_{N-1}(x_{N-1}, d_{N-1}) = \min_{\substack{x_N \in \mathcal{X}, \\ u_{N-1} \in \mathcal{U}}} \; a_{N-1}^T x_{N-1} + b_{N-1}^T u_{N-1} \tag{6.2.4a}$$

$$\text{s.t. } x_N = Ax_{N-1} + Bu_{N-1} + Cd_{N-1} \tag{6.2.4b}$$

The cost-to-go functions $Q_t(\cdot)$ are continuous, piece-wise linear, and convex (Pereira and Pinto, 1991). This allows us to construct a dual dynamic programming (DDP) scheme in which the cost-to-go functions are approximated using collection of hyperplanes (cutting planes). In each iteration $i = 0, 1, \ldots$, the DDP algorithm computes a set of policies $\mathbf{u}_{\mathcal{N}}^i := (u_0^i, u_1^i, \ldots, u_{N-1}^i)$ and $\mathbf{x}_{\mathcal{N}}^i := (x_1^i, x_2^i, \ldots, x_N^i)$, and a set of cuts, one for each stage $t = 0, 1, \ldots, N-1$. An iteration sequence of approximate linear programs (LPs) $\underline{Q}_t^i(\cdot)$ is generated. Each of these LPs $\underline{Q}_t^i(\cdot)$ provides a lower bounding function for the corresponding cost-to-go $Q_t(\cdot)$. The approximation $\underline{Q}_t^i(\cdot)$ is given by the LP (6.2.5), which is solved in a forward sweep to update the primal variables (states and controls) x_{t+1}^i, u_t^i, and in a

backward sweep to update the dual variables (adjoints) π_t^i:

$$\underline{Q}_t^i(x_t^i, d_t) = \min_{\substack{x_{t+1}\in\mathcal{X}, u_t\in\mathcal{U}, \\ \theta_{t+1}\in\mathbb{R}}} a_t^T x_t^i + b_t^T u_t + \theta_{t+1} \tag{6.2.5a}$$

$$\text{s.t. } x_{t+1} = Ax_t^i + Bu_t + Cd_t \quad (\pi_t) \tag{6.2.5b}$$

$$\theta_{t+1} \geq v_{t+1}^l + (\pi_{t+1}^l)^T Ax_{t+1},\ l = 0, 1, \ldots, i-1. \tag{6.2.5c}$$

The cutting plane coefficients v_t^i are computed from:

$$v_t^i = \underline{Q}_t^i(x_t^i, d_t) - (\pi_t^i)^T Ax_t^i, \tag{6.2.6}$$

where $\underline{Q}_t^i(x_t^i, d_t)$ is the objective value of the LP (6.2.5) solved in the backward sweep at iteration i with fixed state x_t^i (obtained from the forward sweep in the same iteration i).

LP duality provides basic properties that justify the choice of the cut coefficients v_t^i. If π_t is the dual solution corresponding to the state dynamics, the cost-to-go function can also be represented in dual form as $Q_t(x_t, d_t) = (\pi_t)^T Ax_t + (\pi_t)^T Cd_t$. A coefficient v_t is used to approximate $(\pi_t)^T Cd_t$, so that $v_t = Q_t(x_t, d_t) - (\pi_t)^T Ax_t$. In each iteration i, the objective $\underline{Q}_t^i(x_t^i, d_t)$ and the dual solution π_t^i obtained at stage t in the backward sweep are used to define the cut for the problem corresponding to the previous stage $t-1$, which provides a lower bound for the cost-to-go $\underline{Q}_t^i(x_t^i, d_t)$ that appears in the objective at stage $t-1$. From duality, one can also establish the key property that the cost-to-go functions are supported by a finite number of dual vertices.

In a forward sweep, the stage subproblem minimizes the current stage cost ($a_t^T x_t^i + b_t^T u_t$) and the cost-to-go from next stage (approximated by the cutting planes). The state variable x_t at time t encodes information of the history of the disturbances and control actions implemented in the previous stages. In particular, the cost variable θ_{t+1} encodes information of the future stages. The problem at the first stage provides a lower bound for the objective value of the entire problem, as it includes a lower approximation of the cost-to-go for all the stages. An upper bound can be obtained from the average over all

iterations of the stage costs at the end of each forward sweep. One can guarantee that this indeed provides an upper bound because the policy x_t^i, u_t^i obtained from the forward sweep is feasible but not optimal. From this observation, we can also see that a forward sweep can be interpreted as a receding horizon MPC scheme that computes an approximate policy of the long-horizon problem. In the backward sweep in each iteration, cutting planes are collected and added to the approximate problems of each stage which incrementally bring the objective function closer to the optimal. The duals used to construct the cutting planes can be interpreted as adjoint variables, which propagate information backward in time. The algorithm is stopped when the upper and lower bounds are within a pre-defined tolerance ϵ. Finite convergence is guaranteed because the cost-to-go functions are supported by a finite number of vertices.

The DDP scheme can be summarized as below:

1. *START with* $x_0^0 = x_0$, $v_{t+1}^0 = 0$, $\pi_{t+1}^0 = 0$ *for* $t = 0, 1, \ldots, N-1$, *and set iteration counter* $i \leftarrow 0$.

   Forward Pass:

2. *Compute* $\underline{Q}_0^i(x_0^i, d_0)$, *and get* x_1^i, u_0^i. *Set the current lower bound,* $lb = \underline{Q}_0^i(x_0^i, d_0)$.

3. *For* $t = 1, 2, \ldots, N-1$, *solve* $\underline{Q}_t^i(x_t^i, d_t)$ *and get the policy* x_{t+1}^i, u_t^i.

4. *Set the current upper bound,*
   $ub = \frac{1}{i+1} \sum_{l=0}^{i} \sum_{k=0}^{N-1} a_k^T x_k^l + b_k^T u_k^l$.

   Backward Pass: *For* $t = N-1, N-2, \ldots, 1$,

5. *Compute* $\underline{Q}_t^i(x_t^i, d_t)$ *and obtain associated dual* π_t^i.

6. *Compute the cut coefficient:*
   $v_t^i = \underline{Q}_t^i(x_t^i, d_t) - (\pi_t^i)^T A x_t^i$.
   Add a cut to the problem $\underline{Q}_{t-1}^i(x_{t-1}^i, d_{t-1})$: $\theta_t \geq v_t^i - (\pi_t^i)^T A x_t$.

7. *UPDATE* $i \leftarrow i+1$. *If* $ub - lb \geq \epsilon$, *RETURN to Step 2, else STOP.*

A key advantage of DDP is that it decomposes the time horizon into stages (in this case each stage is a time step) and only solves single-stage optimization problems. This enables the solution of problems with extremely long time horizons. In other words, the algorithm marches sequentially in time and never needs to form the entire problem explicitly (and thus does not need to store it in memory). All information is encoded (summarized) in the form of cutting planes that are used to approximate the cost-to-go functions. Also, since a standard MPC control policy is equivalent to a single sweep of the DDP scheme, we conclude that the performance of standard MPC is suboptimal compared to that of DDP.

6.3 Stochastic Multistage Formulation

We now consider a stochastic MPC formulation that explicitly takes into account the full uncertainty description of the *r*andom process $D_{\mathcal{N}}$. The joint distribution of the process is $p(D_{\mathcal{N}}) = p(D_0, D_1, \ldots, D_{N-1})$ and the conditional distribution of D_t given information up to time $t-1$ is $p(D_t|I_{t-1})$. The conditional expectation of the cost-to-go $Q(x_t, D_t)$ is denoted as $\mathbb{E}_{D_t|I_{t-1}}[Q(x_t, D_t)]$. A key observation is that, since the state contains all necessary information about the past disturbances, the disturbance realizations can be considered independent, and the conditional expectation becomes an expectation with respect to the next disturbance only. We use this notation to formulate a stochastic MPC problem as a multistage SP of the form:

$$Q_0(x_0, d_0) = \min_{\substack{x_1 \in \mathcal{X}, \\ u_0 \in \mathcal{U}}} a_0^T x_0 + b_0^T u_0 + \mathbb{E}_{D_1|I_0}[Q_1(x_1, D_1)] \tag{6.3.7a}$$

$$\text{s.t. } x_1 = Ax_0 + Bu_0 + Cd_0 \tag{6.3.7b}$$

where,

$$Q_1(x_1, d_1) = \min_{\substack{x_2 \in \mathcal{X}, \\ u_1 \in \mathcal{U}}} a_1^T x_1 + b_1^T u_1 + \mathbb{E}_{D_2|I_1}[Q_2(x_2, D_2)] \tag{6.3.8a}$$

$$\text{s.t. } x_2 = Ax_1 + Bu_1 + Cd_1. \tag{6.3.8b}$$

Following the recursion we obtain:

$$Q_{N-1}(x_{N-1}, d_{N-1}) = \min_{\substack{x_N \in \mathcal{X}, \\ u_{N-1} \in \mathcal{U}}} a_{N-1}^T x_{N-1} + b_{N-1}^T u_{N-1} \tag{6.3.9a}$$

$$\text{s.t. } x_N = Ax_{N-1} + Bu_{N-1} + Cd_{N-1} \tag{6.3.9b}$$

In this formulation, the cost-to-go function $Q_t(\cdot)$ is a random variable and thus we use the expected cost-to-go $\mathbb{E}_{D_t|I_{t-1}}[Q_t(x_t, D_t)]$ as measure. The conditional expectations are often approximated by using a samples from the conditional distributions. This sample average approximation (SAA) scheme generates a scenario tree with q scenarios per stage. The total number of nodes in the tree grows exponentially as q^N. As a result, the extensive form of the SAA problem will have a number of state and control variables that grows exponentially. This makes it impossible to form and solve extensive forms of practical multistage SPs.

When the realizations are independent, one can use a compact scenario tree in which the number of nodes scales linearly with N, as $q \cdot N$, as shown in Figure 6.1. This compact tree representation encodes information on all possible paths for the random disturbances. In this representation, a node corresponding to scenario $j = 1, \ldots, q$ at stage $t = 0, \ldots, N$ is denoted by the index (j, t), and the disturbance realization at this node is denoted as $d_{j,t}$. The probability associated with the outcome in node j, t is $p_{j,t}$. We also note that the first stage ($t = 0$) only has one scenario because the disturbance outcome is known (this node is often known as the root node).

6.3.1 *Stochastic Dual Dynamic Programming*

In the stochastic dual dynamic programming (SDDP) scheme, the expected cost-to-go functions $\mathbb{E}_{D_t|I_{t-1}}[Q_t(\cdot)]$ are also lower-approximated using *c*utting planes. This is possible because the expected cost-to-go is still piece-wise linear continuous. SDDP is implemented in the same way as its deterministic counterpart. However, because the disturbance can take multiple possible paths at each stage, a forward sweep is performed on a sample path. A sample path can be represented by $\{(j_t, t)|j_t \in \{1, 2, \ldots, q\}\}$, $t = 0, 1, \ldots, N-1$. In each iteration i, we obtain policies $\mathbf{u}^i_{\mathcal{N}} := (u^i_0, u^i_1, \ldots, u^i_{N-1})$ and $\mathbf{x}^i_{\mathcal{N}} := (x^i_1, x^i_2, \ldots, x^i_N)$ for the sample path in the forward sweep. A set of cuts, one for each stage $t = 0, 1, \ldots, N-1$ (all scenario $j = 1, 2, \ldots, q$ at stage t have the same cut) is obtained in the backward sweep. Consequently, the stage problems in the sample path explored in the forward sweep retain the information of all scenarios at all stages. A sequence of approximate LPs $\underline{Q}^i_{j,t}(\cdot)$ is generated for each stage as a result of the iterations. Each of these LPs provides a lower bound for the corresponding cost-to-go $Q_{j,t}(\cdot)$. The approximate LP $\underline{Q}^i_{j,t}(\cdot)$ is given by the following problem, which is solved to obtain the primal x^i_{t+1}, u^i_t and dual variables $\pi^i_{j,t}$:

$$\underline{Q}^i_{j,t}(x^i_t, d_{j,t}) = \min_{\substack{x_{t+1}\in\mathcal{X}, u_t\in\mathcal{U},\\ \theta_{t+1}\in\mathbb{R}}} a_t^T x^i_t + b_t^T u_t + \theta_{t+1} \tag{6.3.10a}$$

$$\text{s.t. } x_{t+1} = Ax^i_t + Bu_t + Cd_{j,t} \quad (\pi_{j,t}) \tag{6.3.10b}$$

$$\theta_{t+1} \geq \sum_{j=1}^{q} p_{j,t}(v^l_{j,t+1} + (\pi^l_{j,t+1})^T A x_{t+1}),\ l = 0, 1, \ldots, i-1 \tag{6.3.10c}$$

where the cut coefficients $v^i_{j,t}$ are computed from

$$v^i_{j,t} = \underline{Q}^{i+1}_{j,t}(x^i_t, d_{j,t}) - (\pi^i_{j,t})^T A x^i_t \tag{6.3.11}$$

The problem at the first stage can be shown to also provide a lower bound of the optimal objective value of the original stochastic MPC problem. This is because the supporting hyperplanes explored along the sample paths are a subset of the entire set. A statistical

(asymptotically accurate) upper bound for the problem is obtained from the average costs over all iterations at the end of each forward sweep. A history of the α-confidence intervals of the statistical upper bound can be computed by using information from multiple paths. The SDDP scheme can be summarized as:

1. *START with* $x_0^0 = x_0$, $v_{j,t+1}^0 = 0$, $\pi_{j,t+1}^0 = 0$ *for* $j = 1, 2, \ldots, q$, $t = 0, 1, \ldots, N-1$, *and set iteration counter* $i \leftarrow 0$.

   Forward Pass:

2. *Randomly sample a path* $\{(j_t, t) | j_t \in \{1, 2, \ldots, q\}\}$, $t = 0, 1, \ldots, N-1$.

3. *Compute* $\underline{Q}_0^i(x_0^i, d_0)$ *and get policy* x_1^i, u_0^i. *Set the current lower bound,* $lb = \underline{Q}_0^i(x_0^i, d_0)$.

4. *For* $t = 1, 2, \ldots, N-1$, *compute* $\underline{Q}_{j_t,t}^i(x_t^i, d_{j_t,t})$ *and policies* x_{t+1}^i, u_t^i.

5. *Set the current upper bound* $ub = \frac{1}{i+1} \sum_{l=0}^{i} \sum_{k=0}^{N-1} a_k^T x_k^l + b_k^T u_k^l$. *Compute the 95% confidence interval for the statistical upper bound* ub.

   Backward Pass: *For* $t = N-1, N-2, \ldots, 1$ *and for all* $j = 1, 2, \ldots, q$,

6. *Evaluate* $\underline{Q}_{j,t}^i(x_{j,t}^i, d_{j,t})$ *and dual solution* $\pi_{j,t}^i$.

7. *Compute the cut coefficient:* $v_{j,t}^i = \underline{Q}_{j,t}(x_{j,t}^i, d_{j,t}) - (\pi_{j,t}^i)^T A x_t^i$. *Add cut* $\theta_t \geq \sum_{j=1}^{q} p_{j,t-1}(v_{j,t}^i + (\pi_{j,t}^i)^T A x_t)$ *to* (6.3.10).

8. *UPDATE* $i \leftarrow i+1$. *If* $ub - lb \geq \epsilon$, *RETURN to Step 2, else STOP.*

The SDDP scheme can be shown to converge in a finite number of iterations with probability one (Shapiro, 2011; Philpott and Guan, 2008).

6.3.2 *Extended MPC Formulations*

The established connections between MPC and SDDP also reveal possibilities to design MPC schemes that can handle multiple timescales. For instance, the SDDP scheme introduced assumes a single time step per stage, and implicitly assumes that information of the random process is collected after each timestep. When multiple timescales are present in the problem, uncertainty might reveal at a slower timescale than the system dynamics. For instance, in energy systems, market realizations might reveal after each hour while the controller makes decisions at a time resolution of 5 minutes. Such cases lead to internal time steps within the stages. Here, we present modifications needed to provide statistically consistent dual information and guarantee convergence.

Consider the stochastic MPC problem in which the realizations of the disturbances are observed at times t and $t+1$, but the system dynamics evolve over inner stage times $t = t_0$, $t_1, \ldots, t_n$, $t_{n+1} = t+1$. When the realization is observed at t, a trajectory of the disturbance at the internal times $\boldsymbol{d}_{j,t} := (d_{j,t}, d_{j,t_1}, d_{j,t_2}, \ldots, d_{j,t_n})$ is also observed. We also denote the state and control trajectories $\boldsymbol{x}_t := (x_t, x_{t_1}, x_{t_2}, \ldots, x_{t_n})$ and $\boldsymbol{u}_t := (u_t, u_{t_1}, u_{t_2}, \ldots, u_{t_n})$. The cost-to-go for stage t and scenario j is given by

$$Q_{j,t}(\boldsymbol{x}_t, \boldsymbol{d}_{j,t}) = \min_{\substack{x_{t+1} \in \mathcal{X}, \\ u_t \in \mathcal{U}}} \boldsymbol{a}_t^T \boldsymbol{x}_t + \boldsymbol{b}_t^T \boldsymbol{u}_t + \mathbb{E}_{D_{t+1}|I_t}[Q_{t+1}(\boldsymbol{x}_{t+1}, \boldsymbol{D}_{t+1})] \tag{6.3.12a}$$

$$\text{s.t. } x_{t+1} = Ax_{t_n} + Bu_{t_n} + Cd_{j,t_n} \quad (\pi_{j,t_n}) \tag{6.3.12b}$$

$$x_{t+1_{p+1}} = Ax_{t+1_p} + Bu_{t+1_p} + Cd_{j,t+1_p} \quad (\pi_{j,t+1_p}),\ p = 0, 1, \ldots, n-1. \tag{6.3.12c}$$

From duality, we have that the cost-to-go satisfies:

$$Q_{j,t}(\boldsymbol{x}_t, \boldsymbol{d}_{j,t}) = (\pi_{j,t_n})^T A x_{t_n} + \sum_{p=0}^{n-1} (\pi_{t+1_p})^T A x_{t+1_p} + (\pi_{j,t_n})^T C d_{t_n} + \sum_{p=0}^{n-1} (\pi_{t+1_p})^T C d_{t+1_p} \tag{6.3.13}$$

We propose to use the coefficient $v_{j,t}$ (similar to that in (6.3.11)) to approximate the terms not involving the states on the right-hand side of (6.3.13). With this, we have that

$$v_{j,t} = Q_{j,t}(\boldsymbol{x}_t, \boldsymbol{d}_{j,t}) - \left((\pi_{j,t_n})^T A x_{t_n} + \sum_{p=0}^{n-1} (\pi_{t+1_p})^T A x_{t+1_p} \right). \tag{6.3.14}$$

To compute $v^i_{j,t}$ using $x^i_{t_n}$ we use:

$$v^i_{j,t} = Q^i_{j,t}(\boldsymbol{x}^i_t, \boldsymbol{d}_{j,t}) - \left((\pi_{j,t_n})^T A x_{t_n} + \sum_{p=0}^{n-1} (\pi_{t+1_p})^T A x_{t_n} \right). \tag{6.3.15}$$

In the case studies that follow, we demonstrate that this modification ensures convergence. A rigorous convergence analysis will be established in future work.

SDDP can also be used to construct hierarchical MPC schemes. For instance, one can envision a long-term MPC formulation that uses an approximate solution obtained with SDDP to obtain cutting planes that guide a short-term MPC controller operating at high time resolution. In particular, the cutting planes act as a terminal cost that restricts the searchable state and control space of the short-term MPC controller.

6.4 Case Study

We illustrate the scalability of the proposed SDDP scheme using a case study for an electrochemical energy storage system that modulates the load of a building and participates in multiscale electricity markets. A diverse set of energy systems such as generators, batteries, wind turbines, and flywheels can participate in electricity markets. The markets are structured at multiple time levels, namely day-ahead (hourly market commitments) and real-time markets (commitments ranging from minutes to seconds) (Dowling et al., 2017). In day-ahead markets (DAM) the electricity is traded in intervals of 1 hour with the prices being constant in each interval of 1 hour and varying with intervals. Real-time markets (RTM), on the other hand, can have varying timescales depending on the ISO (Indepen-

dent System Operators) operating it. For example, in California markets (CAISO) the real-time markets are operated at two timescales, a 15-minute timescale which is commonly called as the quarter-hourly market (QHM) and a 5-minute timescale which is called as the real-time market (RTM), although both are categorized under the real-time market operations by CAISO. The real-time market is a spot market in which utilities can buy power to meet the last few increments of demand not covered in their day-ahead schedules. The frequency of energy price variation is different for day-ahead and real-time markets (e.g. highly volatile and varying with high frequency in the RTM, while less volatile and varying with low frequency in DAM). A challenge that arises in our context is that capturing these revenue streams at multiple timescales requires the formulation of large MPC problems with long horizons and fine time resolutions.

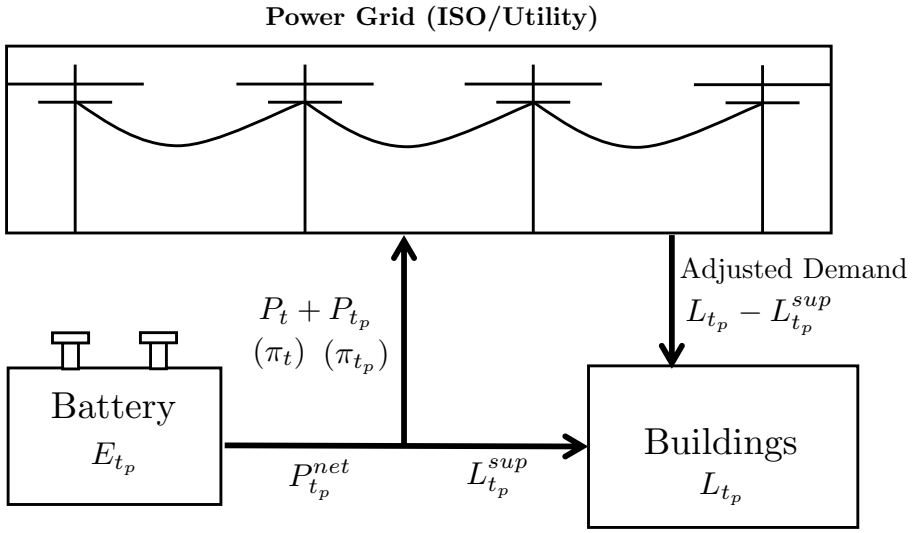

Figure 6.2: Illustration of battery system

We consider a Li-ion battery with an attached building load that is inter-connected to the power grid for providing different services. Batteries can provide power to or draw power from the grid and get compensated at the time-varying prices. The goal is to determine the optimal charge-discharge profile to maximize the revenues generated by providing services to the grid and at the same time meeting the load demands from the building. We consider battery and building as one system (building-battery system) as shown in

Figure 6.2, and any unmet load demand from the building is penalized with the corresponding electricity price in the real-time market. We consider the participation of the battery in the DAM (1-hour timescale) and RTM (5-minute timescale) for a planning horizon of 7 days (168 hours). This problem has a total of 2,016 time steps of 5 minutes. We consider that every stage contains one hour of operation (we have 12 time steps and 168 stages).

To generate an uncertainty characterization for the building loads, we partition real historical data for the loads of a typical building for 1 year (available over 5-minute intervals) into 52 subsets of weekly load profiles (i.e. 2016 intervals of 5-minute each). This division of data helps capture the different load profiles over the weekdays and weekends. We model the weekly load profile as a random process described by a multivariate normal random variable. We obtain a covariance approximation from the historical data by using a Ledoit-Wolf estimator (Kumar et al., 2018a). We generate 50 load scenarios for a week (Figure 6.3) and we use these to generate sample paths.

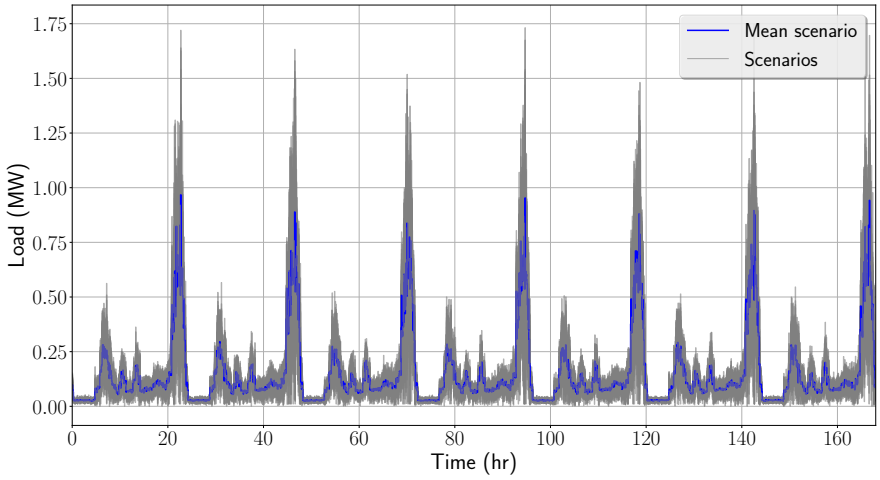

Figure 6.3: Load scenarios for a week of operation.

For the problem for a planning horizon of 7 days, since there are 168 hours and 12 inner intervals in each hour, we have $N = 168$ and $n = 12$. We use π_t to denote the

electricity price in the DAM at hour t and π_{t_p} to denote the electricity price in the RTM at the p^{th} 5-minute interval in hour t. The prices are adjusted based on the corresponding timescale. We use $E_{max} = 0.5$ to denote the energy storage capacity of the battery (in MWh), $P_{max} = 0.5$ to denote the maximum discharging or charging rate of the battery (in MW), and E_0 to denote the initial condition. The random loads of the buildings are denoted as L_{t_p} denoting that loads affect the dynamics in real-time. We use P_t to denote the power committed in DAM at hour t and P_{t_p} to denote power committed in RTM in p^{th} 5-minute interval in hour t. We use $L_{t_p}^{sup}$ to denote the power supplied to the building by the battery and the energy storage level of the battery is E_{t_p} at the p^{th} 5-minute interval in hour t.

The net discharge (power) from the battery is given by:

$$P_{t_p}^{net} = P_t + P_{t_p} + L_{t_p}^{sup},\ t = 0, 1, \ldots, N-1,\ p = 0, 1, \ldots, n. \tag{6.4.16}$$

The energy storage is given by:

$$E_{t+1} = E_{t_n} - P_{t_n}^{net},\ t = 0, 1, \ldots, N-1 \tag{6.4.17}$$

$$E_{t+1_{p+1}} = E_{t+1_p} - P_{t+1_p}^{net},\ t = 0, 1, \ldots, N-1,\ p = 0, 1, \ldots, n. \tag{6.4.18}$$

The states and controls are bounded for $t = 0, 1, \ldots, N-1$, $p = 0, 1, \ldots, n$ as:

$$0 \leq E_{t_p} \leq E_{max} \tag{6.4.19a}$$

$$-P_{max} \leq P_{t_p} \leq P_{max} \tag{6.4.19b}$$

$$-P_{max} \leq P_t \leq P_{max} \tag{6.4.19c}$$

$$-P_{max} \leq P_{t_p}^{net} \leq P_{max} \tag{6.4.19d}$$

$$0 \leq L_{t_p}^{sup} \leq L_{t_p}. \tag{6.4.19e}$$

The objective is given by the total battery revenue:

$$-\sum_{t=0}^{N-1}\sum_{p=0}^{n}\pi_{t_p}P_{t_p}-\sum_{t=0}^{N-1}\pi_t P_t+\sum_{t=0}^{N-1}\sum_{p=0}^{n}\pi_{t_p}(L_{t_p}-L_{t_p}^{sup}). \tag{6.4.20}$$

Given the number of stages $N = 168$ and $q = 50$ outcomes per stage, the extensive form of the associated multi-stage SP would involve an scenario tree with $50^{168} = 2.67 \times 10^{285}$ nodes, which is clearly intractable. SDDP overcomes this scalability bottleneck by sampling paths on-the-fly.

The SDDP algorithm is implemented in `Julia` programming language using `JuMP` algebraic modeling package (Dunning et al., 2017). We solve all stage subproblems in the SDDP scheme using `Gurobi 7.0.1`. For a typical run of the SDDP algorithm with $\epsilon = 0.1$, we obtain the following performance numbers:

- Number of SDDP iterations: 221
- Time to solution: 825.8 sec
- Final lower bound on cost: -554.88
- Final upper bound on cost: -554.79 with a confidence interval of [-554.57, -555.02]

A typical trajectory obtained for the lower and the upper bound of the cost is presented in Figure 6.4. As the number of iterations increases the upper bound comes closer to the lower bound and its confidence interval shrinks as more sample path information is accumulated. This convergence behavior shows that the proposed multiscale SDDP variant performs satisfactorily.

The optimal revenue obtained from SDDP is about \$554.84. To provide a comparison, a two-stage relaxation of the multi-stage problem provides an optimal revenue of \$711.08. The policy obtained with this relaxation is not feasible but provides an idea of the value of rigorously capturing uncertainty behavior.

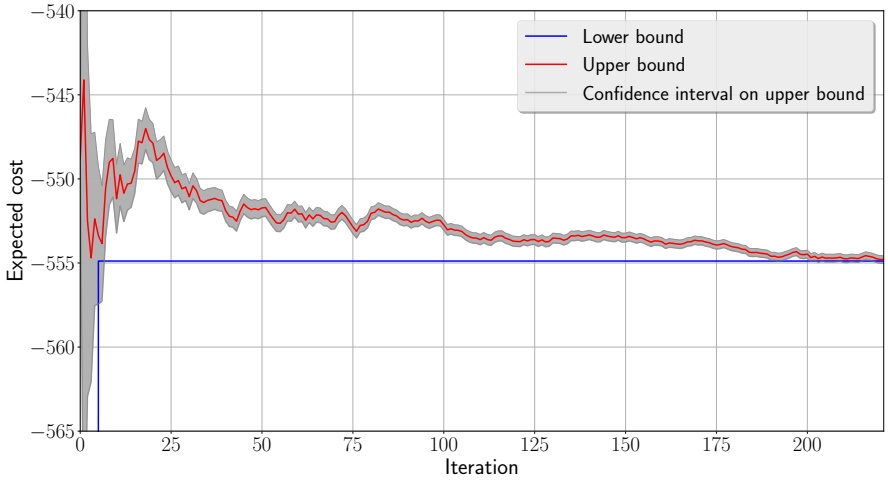

Figure 6.4: Evolution of lower and upper bounds in SDDP.

6.5 Conclusions

In this chapter, we demonstrated that stochastic dual dynamic programming (SDDP) provides a scalable framework to address complex MPC applications that need to capture long horizons and fine time resolutions. By deriving and interpreting SDDP from an MPC perspective, we showed that important insights can be gained on how to construct terminal costs and performance bounds for MPC. In Chapter 7, we leverage the nested cutting plane scheme with forward and backward sweeps along the time horizon in the SDDP algorithm for mixed-integer MPC problems that include discrete or binary control decisions and/or states.

7

DUAL DYNAMIC PROGRAMMING FOR MULTISCALE MIXED-INTEGER MPC

In this chapter, we extend the nested cutting plane scheme of SDDP, which uses forward and backward sweeps along the time horizon to adaptively construct and refine the cost-to-go functions to develop a dual dynamic programming approach for solving multiscale mixed-integer model predictive control (MPC) problems that can handle real constraints of energy systems such as discrete actuators and minimum startup loads.

7.1 Introduction

Model predictive control (MPC) is a well-known control technology for a wide range of industrial processes due to its capability to handle any complex dynamic models with a variety of control actuators and system constraints. However, integer or binary (discrete) actuators that are common in many large-scale industrial applications and energy systems have not been well studied in the literature of MPC. The major issue with discrete controls is that they result in mixed-integer optimization problems which makes the problem computationally expensive to solve (García et al., 1989; Mayne et al., 2000; Rawlings et al., 2017; Kirches, 2011). To handle the discrete controls in the industrial applications of MPC, they

are dealt with separately outside the MPC control layer and those discrete decisions are determined using heuristics or other logical rules in a different automation layer. However, with advances in the optimization algorithms for mixed-integer linear programs (MILP), recently MPC with mixed-integer models have been studied and the theory of MPC has been extended to include discrete control actions (Rawlings and Risbeck, 2017). Various results on stability of mixed-integer MPC have also appeared in the literature (Bemporad and Morari, 1999; Di Cairano et al., 2014; Aguilera and Quevedo, 2013; Rawlings and Risbeck, 2017).

Recent works have shown that it is very important to represent systems with both continuous and discrete controls using mixed-integer models (Bemporad and Morari, 1999; Camacho et al., 2010; Kobayashi et al., 2014). Large-scale energy systems such as central HVAC plants for campuses or districts have both continuous and discrete controls decisions to make in real-time (Rawlings et al., 2018; Risbeck et al., 2015). In particular, for real-time applications, it is very important to solve large-scale mixed-integer MPC problems arising in these systems efficiently and scalably. Long horizons and fine time discretizations are often required to properly account for long-term effects such as demand charges and short-term disturbances at high frequencies such as energy loads and prices (Kumar et al., 2018a, 2019a, 2020). Large-scale mixed-integer problems are also commonly encountered in integrated production scheduling and control for industrial applications (Daoutidis et al., 2018; Baldea et al., 2015; Beal et al., 2017, 2018). Recent work on online schedule or rescheduling in chemical production applications also involves dealing with large-scale mixed-integer models that are computationally challenging and need to be solved faster (Subramanian et al., 2012; Gupta and Maravelias, 2017, 2016; Pattison et al., 2017; Gupta et al., 2016).

Large-scale mixed-integer problems arising in these applications require efficient solution algorithms for fast solutions. Many solution schemes based on temporal decomposition have been developed for production scheduling problems (Bassett et al., 1996; Harjunkoski and Grossmann, 2001; Jackson and Grossmann, 2003), unit commitment in

power systems (Kim et al., 2018), and hydrothermal scheduling (Sifuentes and Vargas, 2007). We highlight that the decomposition schemes presented in these works are based on Benders decomposition, and are limited by the size of the master problem, and are often not scalable. In (Pereira, 1989) and (Pereira and Pinto, 1991), the authors present a scalable algorithm, namely stochastic dual dynamic programming (SDDP), based on a nested Benders decomposition scheme that solves small problems spanning a single time-step in forward and backward sweeps along the time horizon while adding cutting planes in each iteration. However, the SDDP algorithm can be applied to large-scale continuous linear programs arising from multi-stage stochastic programming problems. Recently, the SDDP scheme was derived and interpreted for long-horizon multiscale stochastic MPC problems and it was implemented for electricity market participation of batteries (Kumar et al., 2018c). Recent works also explore the use of nested Benders decomposition scheme with forward and backward sweeps of SDDP framework for multi-stage stochastic programs with binary state variables (Zou et al., 2016, 2019), and the algorithm named as stochastic dual dynamic integer programming (SDDIP). It is important to note that the SDDIP schemes presented in (Zou et al., 2016, 2019) are limited to systems with continuous control decisions. The authors in (Lara et al., 2018) and (Lara et al., 2019) have also implemented the SDDIP algorithm for both deterministic and stochastic electric power infrastructure planning which allows continuous and integer state variables. The implementations of the SDDIP schemes, however, need to be derived in the context of mixed-integer MPC applications where both controls and states can be continuous and integer variables.

In this work, we interpret and derive a dual dynamic integer programming (DDIP) scheme for solving multiscale mixed-integer model predictive control (MPC) problems that can handle real constraints of energy systems such as discrete controls and minimum startup loads. We leverage the nested cutting plane scheme used in the stochastic dual dynamic programming for multi-stage stochastic programming problems, which uses forward and backward sweeps along the time horizon to adaptively construct and refine the cost-to-go functions. We highlight that the DDIP scheme presented in this work

can handle both continuous and discrete (or binary) control and state variables in the mixed-integer MPC problems. This approach provides a scalable algorithm to solve the large-scale mixed-integer programs arising in a wide range of applications such as MPC with long horizons or fine discretizations, online scheduling, and integrated production scheduling and control.

We demonstrate the performance of the proposed dual dynamic integer programming scheme through computational experiments for a mixed-integer MPC problem for a central HVAC plant for a typical university campus. We show that the proposed scheme is able to obtain the optimal solution within an optimality gap of 0.1% for the LP relaxation (continuous form) of the MPC problem for the central HVAC plant. Then, we present the results for the implementation of the dual dynamic integer programming scheme to the mixed-integer MPC problem for the central HVAC plant case with different horizon lengths and discretizations to show that the DDIP scheme provides a scalable approach to solve large-scale mixed-integer MPC problems.

7.2 Mixed-Integer MPC Formulation

We consider a mixed-integer MPC problem for a linear system of the following form:

$$\min_{\mathbf{x}_{\mathcal{N}}, \mathbf{u}_{\mathcal{N}}} \sum_{t=0}^{N-1} f_t(x_t, u_t) \tag{7.2.1a}$$

$$\text{s.t. } x_{t+1} = Ax_t + Bu_t,\ t = 0, \ldots, N-1 \tag{7.2.1b}$$

$$x_0 = \bar{x} \tag{7.2.1c}$$

$$u_t \in \mathcal{U}, u_t \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}},\ t = 0, \ldots, N-1 \tag{7.2.1d}$$

$$x_t \in \mathcal{X}, x_t \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}},\ t = 0, \ldots, N. \tag{7.2.1e}$$

Here, the controls and states at time t are denoted as u_t and x_t, respectively. $f_t(x_t, u_t)$ is the stage cost at time t associated with the states x_t and controls u_t. The trajectory of control u_t is defined over a horizon of N time stages from $t = 0$ to $t = N-1$, while the

trajectory of state x_t is defined from $t = 0$ to $t = N$ with the initial state provided as $x_0 = \bar{x}$. We note that the states x_t and controls u_t have n_{u_1} and n_{x_1} integer components, and n_{u_2} and n_{x_2} continuous components, respectively. The states $x_t \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}}$ and controls $u_t \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}}$ (with $n_x = n_{x_1} + n_{x_2}$, $n_u = n_{u_1} + n_{u_2}$) are also bounded by the polyhedral sets $\mathcal{X}$ and $\mathcal{U}$, respectively. The dynamics of the linear system are described by the matrices $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$. In this work, we assume that the stage costs $f_t(x_t, u_t)$ are also linear functions of the controls and states.

7.2.1 *Dynamic Programming Representation*

We now consider a reformulation of the MPC problem (7.2.1) using a dynamic programming representation in a multi-stage form (Bertsekas et al., 1995). In this formulation, the objective function is expressed as the sum of the current stage cost and a cost-to-go function that summarizes the cost of future stages. The MPC problem takes the following form as a dynamic program:

$$Q_0(x_0) = \min_{x_1, u_0} f_0(x_0, u_0) + Q_1(x_1) \tag{7.2.2a}$$

$$\text{s.t. } x_1 = Ax_0 + Bu_0 \tag{7.2.2b}$$

$$u_0 \in \mathcal{U}, u_0 \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.2c}$$

$$x_1 \in \mathcal{X}, x_1 \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}}. \tag{7.2.2d}$$

where the state x_0 is given as the initial state $\bar{x}$, and the cost-to-go $Q_1(x_1)$ is defined as:

$$Q_1(x_1) = \min_{x_2, u_1} f_1(x_1, u_1) + Q_2(x_2) \tag{7.2.3a}$$

$$\text{s.t. } x_2 = Ax_1 + Bu_1 \tag{7.2.3b}$$

$$u_1 \in \mathcal{U}, u_1 \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.3c}$$

$$x_2 \in \mathcal{X}, x_2 \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}}. \tag{7.2.3d}$$

Following the recursion, the cost-to-go for stage t, $Q_t(x_t)$ is obtained as:

$$Q_t(x_t) = \min_{x_{t+1}, u_t} f_t(x_t, u_t) + Q_{t+1}(x_{t+1}) \tag{7.2.4a}$$

$$\text{s.t. } x_{t+1} = Ax_t + Bu_t \tag{7.2.4b}$$

$$u_t \in \mathcal{U}, u_t \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.4c}$$

$$x_{t+1} \in \mathcal{X}, x_{t+1} \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}}. \tag{7.2.4d}$$

The cost-to-go at the final stage of the horizon $N-1$ is obtained as:

$$Q_{N-1}(x_{N-1}) = \min_{x_N, u_{N-1}} f_{N-1}(x_{N-1}, u_{N-1}) \tag{7.2.5a}$$

$$\text{s.t. } x_N = Ax_{N-1} + Bu_{N-1} \tag{7.2.5b}$$

$$u_{N-1} \in \mathcal{U}, u_{N-1} \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.5c}$$

$$x_N \in \mathcal{X}, x_N \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}}. \tag{7.2.5d}$$

The dynamic programming representation of the mixed-integer MPC reveals an opportunity to construct iterative schemes that use approximations of the cost-to-go functions. However, to facilitate the approximation of the cost-to-go function, we further introduce an auxiliary variable $z_t \in \mathbb{R}^{n_x}$ as a copy of x_t in Problem (7.2.4). Note that z_t is relaxed to be a continuous variable within the same bounds as state x_t.

$$Q_t(x_t) = \min_{x_{t+1}, u_t, z_t} f_t(x_t, u_t) + Q_{t+1}(x_{t+1}) \tag{7.2.6a}$$

$$\text{s.t. } x_{t+1} = Az_t + Bu_t \tag{7.2.6b}$$

$$z_t = x_t \quad (\mu_t) \tag{7.2.6c}$$

$$u_t \in \mathcal{U}, u_t \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.6d}$$

$$x_{t+1} \in \mathcal{X}, x_{t+1} \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}} \tag{7.2.6e}$$

$$z_t \in \mathbb{R}^{n_x}. \tag{7.2.6f}$$

We highlight that with the introduction of the auxiliary variable z_t, it becomes the only variable linking the cost-to-functions between stages and the constraint (7.2.6c) becomes the only linking constraint between stages. We also note that the dual variable (adjoint) for the linking constraint (7.2.6c) in the corresponding linear programming (LP) relaxation of Problem (7.2.6) is denoted by μ_t, which will be used to design of the dual dynamic integer programming scheme that is described in Section 7.2.2. We will also see that the introduction of the auxiliary variable z_t facilitates the design of the DDIP scheme for a multiscale mixed-integer MPC in Section 7.2.3.

7.2.2 *Dual Dynamic Integer Programming*

We use the dynamic programming formulation (7.2.6) to design a dual dynamic integer programming (DDIP) scheme in which the cost-to-go functions are approximated using a collection of cutting planes (hyperplanes) accumulated in each iteration. However, the cutting planes can be constructed only if the cost-to-go functions are continuous and convex. So, we obtain the cutting planes in the DDIP scheme from the linear programming (LP) relaxations of the cost-to-go functions. In the DDIP algorithm, each iteration $k = 0, 1, \ldots$ comprises a forward sweep and a backward sweep along the time horizon. A forward sweep in iteration k computes a control trajectory $(u_0^k, u_1^k, \ldots, u_{N-1}^k)$ and state trajectory $(x_1^k, x_2^k, \ldots, x_N^k)$ with the given initial state by solving a sequence of approximate problems $\phi_t^k(\cdot)$ (instead of $Q_t(\cdot)$) for $t = 0, 1, \ldots, N-1$. The approximate cost-to-go $\phi_t^k(\cdot)$ is given by the Problem (7.2.7). In the backward sweep along the time horizon in iteration k, a set of Benders cuts, one for each stage $t = N-1, N-2, \ldots, 0$ is generated by solving the LP relaxation of the Problem (7.2.7). The solution to the LP relaxation of the Problem (7.2.7) gives the objective value denoted by $\hat{\phi}_k^t$ and the dual variable μ_t^k for the linking constraint, both of which are used as the coefficients to construct the Benders cuts. The dual variables μ_t^k and the LP relaxation objectives $\hat{\phi}_t^k$ for the Benders cuts at all times are updated in each

iteration k in the backward sweep.

$$\phi_t^k(x_t^k) = \min_{x_{t+1},u_t,z_t,\theta_{t+1}} f_t(x_t^k, u_t) + \theta_{t+1} \tag{7.2.7a}$$

$$\text{s.t. } x_{t+1} = Az_t + Bu_t \tag{7.2.7b}$$

$$z_t = x_t^k \quad (\mu_t^k) \tag{7.2.7c}$$

$$\theta_{t+1} \geq \phi_{t+1}^l + (\mu_{t+1}^l)^T(x_{t+1} - x_{t+1}^l),\ l = 0, 1, \ldots, k-1. \tag{7.2.7d}$$

$$u_t \in \mathcal{U}, u_t \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}} \tag{7.2.7e}$$

$$x_{t+1} \in \mathcal{X}, x_{t+1} \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}} \tag{7.2.7f}$$

$$z_t \in \mathbb{R}^{n_x}. \tag{7.2.7g}$$

Here, the constraints (7.2.7d) are the collection of Benders cuts for stage t accumulated over all previous iterations. In each iteration k for each stage t, the variable θ_{t+1} is used to approximate the cost-to-go for the next stage. We can justify the choice of the cut coefficients in (7.2.7d) using LP duality properties. For the LP relaxation of Problem (7.2.7), if the dual solution corresponding to the linking constraint (7.2.7c) is μ_t^k, the corresponding cost-to-go can also be represented in dual form as $\hat{\phi}_t^k(x_t^k) = (\mu_t^k)^T x_t^k + v_t^k$, where v_t^k denotes the cost component in dual form corresponding to the constraints defining the polyhedral sets $\mathcal{U}$ and $\mathcal{X}$. We can see that if μ_t^k and v_t^k are substituted using information from the previous iteration by μ_t^{k-1} and $\hat{\phi}_t^{k-1} - (\mu_t^{k-1})^T x_t^{k-1}$, respectively, we obtain the Benders cut for stage $t-1$ for the current iteration k: $\theta_t \geq (\mu_t^{k-1})^T(x_t^k - x_t^{k-1}) + \hat{\phi}_t^{k-1}$. In each iteration k, the objective $\hat{\phi}_t^k(x_t^k)$ and the dual solution μ_t^k obtained at stage t in the backward sweep are used to define the Benders cut for the problem at the previous stage $t-1$ for the next iteration $k+1$, which provides a lower bound for the cost-to-go $\phi_t^k(x_t^k)$ that appears in the objective at stage $t-1$. We can also observe a key property arising from duality that the cost-to-go functions are supported by a finite number of dual vertices.

In a forward sweep, the stage subproblem minimizes the current stage cost ($f_t(x_t, u_t)$) and the cost-to-go from next stage (approximated by the cutting planes). The state vari-

able x_t at time t encodes information of the history of the control actions implemented in the previous stages. In particular, the cost variable θ_{t+1} encodes information of the future stages. We note that in iteration k, solving Problem (7.2.7) in the forward sweep gives a set of trajectories for the controls and states $\{u_t^k, x_{t+1}^k\}, t = 0, \ldots, N-1$, starting with the given initial state $x_0 = \bar{x}$. These control and state policies also satisfy the constraints of the MPC Problem (7.2.1) and therefore, are feasible but not necessarily optimal. Thus, we compute the value of $\sum_{t=0}^{N-1} f_t(x_t^k, u_t^k)$ that is guaranteed to be an upper bound for the MPC objective. We can also interpret a forward sweep of the DDIP scheme as a receding horizon MPC scheme that computes an approximate policy of the long-horizon problem. In the backward sweep in each iteration, cutting planes are collected and added to the approximate problems $\phi_t^k(\cdot)$ of each stage which incrementally bring the objective function closer to the optimal. The duals of the LP relaxation problem which are used to construct the cutting planes can be interpreted as adjoint variables that propagate information backward in time. We note that the objective value $\hat{\phi}_0^k(x_0^k)$ provides a lower bound for the MPC objective $Q_0(x_0)$ because $\hat{\phi}_0^k(x_0^k)$ is obtained from the solution to the LP relaxation of Problem (7.2.7) and also contains lower approximations of cost-to-go functions at all stages. The algorithm is stopped when the upper and lower bounds are within a pre-defined tolerance ϵ.

The DDIP scheme can be summarized as below:

1. *START with* $x_0^0 = \bar{x}$, $\hat{\phi}_{t+1}^0 = 0$, $\mu_{t+1}^0 = 0$ *for* $t = 0, 1, \ldots, N-1$, *and set iteration counter* $k \leftarrow 0$.

   Forward Pass:

2. *Compute* $\phi_0^k(x_0^k)$, *and get* x_1^k, u_0^k.

3. *For* $t = 1, 2, \ldots, N-1$, *solve* $\phi_t^k(x_t^k)$ *and get the policy* x_{t+1}^k, u_t^k.

4. *Set the current upper bound,*
   $ub = \sum_{t=0}^{N-1} f_t(x_t^k, u_t^k)$.

Backward Pass:

5. *For $t = N-1, N-2, \ldots, 1$, solve LP relaxation of the updated Problem (7.2.7), and obtain $\hat{\phi}_t^k(x_t^k)$, x_t^k, and associated dual μ_t^k.*

6. *For $t = N-1, N-2, \ldots, 1$, compute the corresponding Benders cuts: $\theta_t \geq \hat{\phi}_t^k + (\mu_t^k)^T(x_t - x_t^k)$; add the computed cuts to the problems $\phi_{t-1}^{k+1}(x_{t-1}^{k+1})$ for future iterations.*

7. *Set the current lower bound, $lb = \hat{\phi}_0^k(x_0^k)$.*

8. *UPDATE $k \leftarrow k+1$. If $ub - lb \geq \epsilon$, RETURN to Step 2, else STOP.*

A key advantage of DDIP scheme is that the time horizon is decomposed into stages and it only needs to solve single time-step optimization problems as each stage comprises a single time step here. This feature enables the solution of problems with extremely long time horizons even in the mixed-integer MPC problems. In other words, the DDIP algorithm moves sequentially in time and never needs to form the entire mixed-integer problem explicitly (and thus storing it in memory is also avoided). The cutting planes that approximate the cost-to-go functions encode (summarize) all the accumulated information. We also note that in each iteration, the backward sweep of the DDIP scheme solves single-stage problems with fixed initial states provided from the forward sweep in the same iteration, and therefore each stage subproblem is independent of other stage subproblems. Thus, all stage subproblems in the backward sweep can be solved in parallel which can further accelerate the algorithm.

7.2.3 *Multiscale Mixed-integer MPC Formulations*

The established connections between mixed-integer MPC and DDIP also reveal possibilities to design DDIP scheme for MPC that can handle multiple timescales. In the DDIP scheme introduced for the mixed-integer MPC in Section 7.2.2 assumes a single time step per stage, and implicitly assumes that information is collected after each time step. When

multiple timescales are present in the problem, we can design a DDIP scheme by decomposing the time horizon into stages such that each stage also contains inner time steps. This means that the DDIP scheme collects information at a slower timescale (stages) than the system dynamics (that evolves over inner time steps). Such a multiscale design of DDIP scheme solves is required when the MPC problem involves an extremely long time horizon due to fine time resolutions. Here, we present modifications needed to update the dual information for the multiscale MPC problem.

Consider the mixed-integer MPC problem in which the DDIP information (cutting planes) are collected at stages t and $t+1$, but the system dynamics evolve over inner stage times $t = t_0, t_1, \ldots, t_{n-1}, t_n, t_{n+1} = t+1$. At stage t, given the state x_t at the start of the stage (time instant t_0), the cost-to-go $Q_t(x_t)$ is solved to obtain the state and control trajectories $\boldsymbol{x_{t+1}} := (x_{t_1}, x_{t_2}, \ldots, x_{t_n}, x_{t+1})$ and $\boldsymbol{u_t} := (u_t, u_{t_1}, u_{t_2}, \ldots, u_{t_n})$. The cost-to-go $Q_t(x_t)$ for stage t is given by

$$Q_t(x_t) = \min_{\boldsymbol{x_{t+1}}, \boldsymbol{u_t}} f_t(\boldsymbol{x_t}, \boldsymbol{u_t}) + Q_{t+1}(x_{t+1}) \tag{7.2.8a}$$

$$\text{s.t. } x_{t_1} = Az_t + Bu_t \tag{7.2.8b}$$

$$x_{t_{p+1}} = Ax_{t_p} + Bu_{t_p},\ p = 1, \ldots, n. \tag{7.2.8c}$$

$$z_t = x_t \quad (\mu_t) \tag{7.2.8d}$$

$$u_{t_p} \in \mathcal{U}, u_{t_p} \in \mathbb{Z}^{n_{u_1}} \times \mathbb{R}^{n_{u_2}},\ p = 0, \ldots, n. \tag{7.2.8e}$$

$$x_{t_p} \in \mathcal{X}, x_{t_p} \in \mathbb{Z}^{n_{x_1}} \times \mathbb{R}^{n_{x_2}},\ p = 1, \ldots, n. \tag{7.2.8f}$$

$$z_t \in \mathbb{R}^{n_x}. \tag{7.2.8g}$$

From duality, we have that the LP relaxation of the cost-to-go $Q_t(x_t)$ satisfies $\hat{Q}_t(x_t) = (\mu_t)^T x_t + v_t$ similar to that discussed in Section 7.2.2. It is easy to see that the DDIP scheme can be designed for the multiscale mixed-integer MPC problem using a Benders cut obtained for each stage t in the backward sweep in each iteration k. We use the computed LP relaxation objective $\hat{\phi}_t^k$, state $x_{t_1}^k$, and dual μ_t^k in iteration k to define the Benders cut for

stage $t-1$ for future iterations: $\theta_t \geq \hat{\phi}_t^k + (\mu_t^k)^T(x_{t_1} - x_{t_1}^k)$ (similar to that in (7.2.7d)).

We also observe that DDIP can also be used to construct hierarchical MPC schemes. For instance, a long-term mixed-integer MPC formulation can be decomposed into hierarchies of stages comprising coarse to fine time resolutions, and in each hierarchy, a DDIP scheme can be used to obtain cutting planes that guide the MPC controller at the lower hierarchy operating at higher time resolution. In particular, the cutting planes act as a terminal cost that restricts the searchable state and control space of the lower level MPC controller. This hierarchical scheme presents a scalable approach to solve long-term mixed-integer MPC problems and will be an interesting area of future work.

7.3 Computational Experiments

In this section, we present some computational experiments to demonstrate the performance of the proposed DDIP schemes. We describe the decision-making setting and the physical dynamic model of a central HVAC plant for a typical university campus, and then present the results of the implementations of the DDIP scheme. The nomenclature used in this section for the HVAC central plant model is provided in Section A.3 of Appendix A.

7.3.1 *Decision-Making Setting*

The decision-making setting of the central HVAC plant for a typical university campus is an extension to that described in Chapter 4. The HVAC plant that we consider in this work consists of a chiller subplant comprising 4 chillers that produce chilled water and a heat recovery (HR) chiller subplant comprising 3 HR chiller units that produce both chilled water and hot water, 3 hot water generators to produce hot water, 9 cooling towers to reduce the temperature of the water purchased from the market, a dump heat exchanger (dump HX) for rejecting heat from the hot water, and storage tanks (one for chilled water and one for hot water). The goal is to determine hourly operating strategies for each equipment unit to minimize the total cost of the external utilities that need to be purchased from the

market (electricity, water, and natural gas). Electricity is charged based on time-varying prices, while water and natural gas usage are charged at respective constant prices. We do not consider the monthly peak demand charges for electricity in this work. .5, .8, .5, .5, 0 Thus, the various cost components for the central plant are:

- *Electricity transactions (hourly)*: Electricity is required for the equipment operation in the central plant and needs to be purchased. The transactions are charged at the time-varying market price, π_t^e.
- *Water transactions (hourly)*: Water is required to make up for evaporative losses of water in the cooling towers. Water is purchased from the utility at a fixed price of π_t^w= \$0.009/gal.
- *Natural gas transactions (hourly)*: Natural gas is needed to run the hot water generators to satisfy the campus heating load. Natural gas is purchased from the utility at a fixed price of π_t^{ng}= \$0.018/kWh.

Figure 7.1 shows the energy flows between all the units of the central HVAC plant and interactions with campus loads and utilities. Electricity is consumed by the chiller subplant, HR chiller subplant, hot water generators, and cooling towers, while utility water is consumed by only the cooling towers to make up for the evaporative losses, and natural gas is only consumed by the hot water generators. The electricity, water, and natural gas consumption of these units are linked to their operating loads. For instance, all chiller units and HR chiller units consume α_{cs}^e and α_{hrc}^e kW of electricity per kW of chilled water produced, respectively; each hot water generator consumes α_{hwg}^e kW of electricity and α_{hwg}^{ng} kW of natural gas per kW hot water produced, respectively; and each cooling tower consumes α_{ct}^e kW of electricity and α_{ct}^w utility water per kW of condenser water input, respectively. Also, there are prescribed minimum operating loads for the chillers, HR chillers, hot water generators, and cooling towers. For instance, each chiller can either be OFF or operate at 50% or more of its maximum cooling capacity when ON. Similarly, each HR chiller unit, when ON, can only operate at minimum 80% of its maximum cooling capacity, each hot

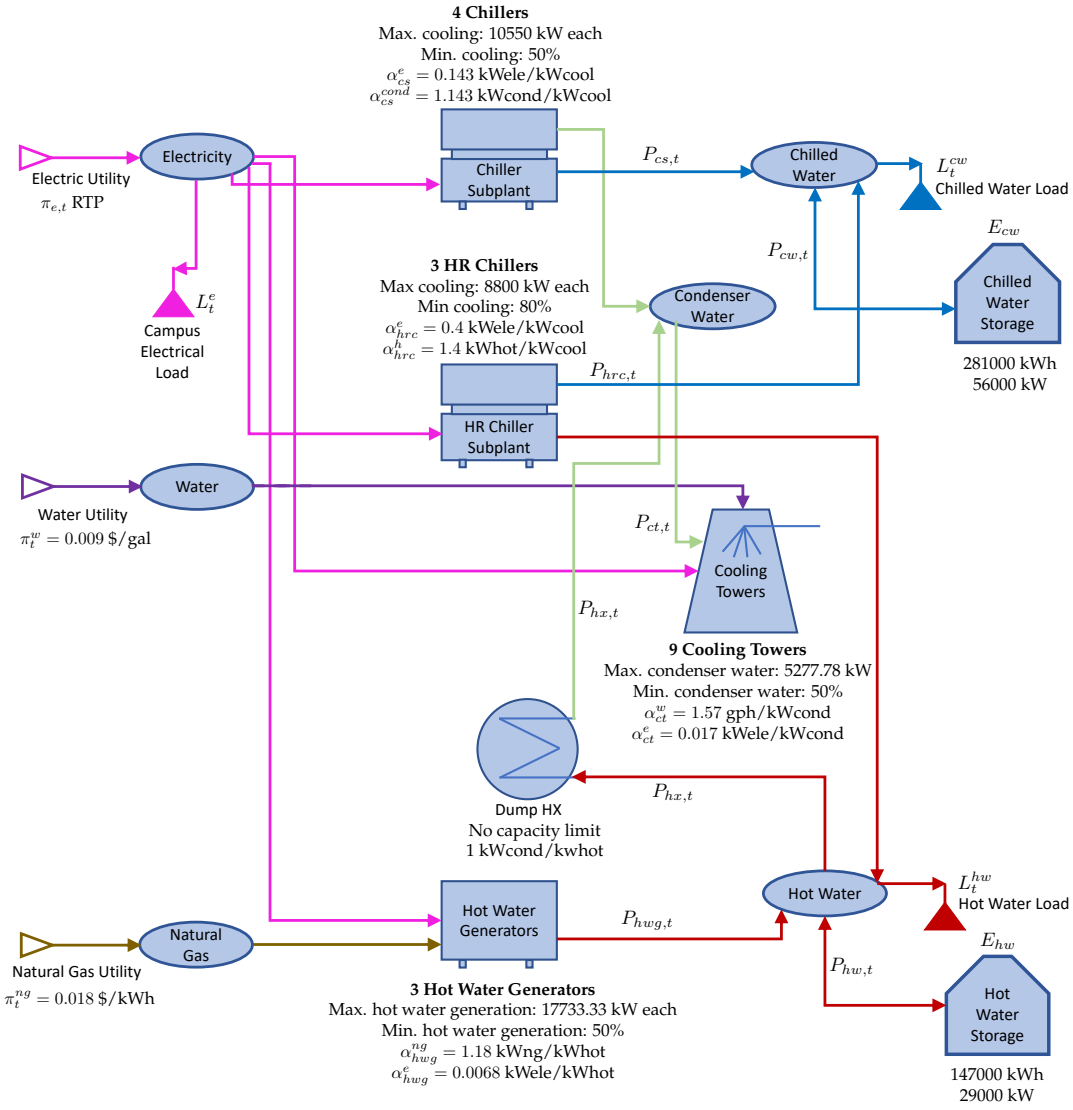

Figure 7.1: Schematic representation of central HVAC plant under study.

water generator and each cooling tower can operate at 50% or more of the respective maximum capacities, while the dump HX can operate at any load when ON. These restrictions require enforcing logical constraints and making binary decisions for each of these units whether to keep a unit ON or OFF, thus resulting in a mixed-integer formulation for the central plant operations.

The central plant meets the chilled water load (L^{cw}_t) of the campus by producing chilled water from the chiller subplant ($P_{cs,t}$), the HR chiller subplant ($P_{hrc,t}$), and the discharge

from chilled water storage ($P_{cw,t}$), where, within the chiller subplant, n^{th} unit produces $p_{n,cs,t}$ chilled water, and within the HR chiller subplant, n^{th} unit produces $p_{n,hrc,t}$ chilled water. The hot water production from the HR chiller subplant ($\alpha^h_{hrc} P_{hrc,t}$), the hot water generators ($P_{hwg,t}$), and the discharge from the hot water storage ($P_{hw,t}$) meets the hot water load (L^{hw}_t) of the campus, where n^{th} HR chiller unit produces $\alpha^h_{hrc} p_{n,hrc,t}$ hot water, and n^{th} hot water generator unit produces $p_{n,hwg,t}$ hot water. The excess hot water ($P_{hx,t}$) in the system is recycled by dump heat exchanger (HX) by cooling it and producing condenser water which is cooled further by the cooling towers together with the condenser water produced by the chiller and HR chiller subplants (total $P_{ct,t}$ condenser water is cooled by the towers, where $p_{n,ct,t}$ is cooled by n^{th} tower). The binary control decisions for the system include the ON/OFF decisions for each of the chiller units, HR chiller units, hot water generators, and cooling towers. The continuous control decisions are the operating loads of all units, which include the chilled water production by each chiller and HR chiller unit (when ON), hot water production by each hot water generator (when ON), the cooling load of each cooling tower (when ON), the heat exchange load of the dump HX, and discharge rates from the two storage tanks.

The HVAC plant operations are driven by uncertain and time-varying disturbances, which are given by the campus loads for electricity (L^e_t), chilled water (L^{cw}_t), and hot water (L^{hw}_t), and by the electricity prices (π^e_t). In this work, we assume that the forecasts for each of these disturbances are available for use in the deterministic MPC formulation. The goal of the control (management) system of the plant is to determine operating loads for all units and storage levels to meet campus loads and to minimize the overall plant cost for the given forecasts. We use similar notation to that in Chapter 4 to describe the mixed-integer MPC problem for the central plant. Section A.3 of Appendix A provides a detailed description of all the variables and quantities involved in this formulation.

7.3.2 *Mixed-Integer MPC Formulation for Central HVAC Plant*

The mixed-integer MPC controller uses the disturbance forecasts for L_t^e, L_t^{cw}, L_t^{hw}, and π_t^e to find the control policy that minimizes the total cost over an N-hour prediction horizon, with the time sets $\mathcal{T} := \{0, \dots, N\}$ and $\bar{\mathcal{T}} := \mathcal{T} \setminus \{N\}$, by solving the following optimization problem:

$$\min \sum_{t \in \bar{\mathcal{T}}} \left(\sum_{j=\{e,w,ng\}} \pi_t^j r_t^j + \sum_{j \in \{cw,hw\}} \rho_j U_{j,t} \right). \tag{7.3.9a}$$

$$\text{s.t. } Y_{n,j,t} \in \{0,1\},\ j \in \{cs, hrc, hwg, ct, hx\}, n \in \{1, \dots, N_j\}, t \in \bar{\mathcal{T}} \tag{7.3.9b}$$

$$\beta_j \overline{p}_j Y_{n,j,t} \le p_{n,j,t} \le \overline{p}_j Y_{n,j,t},\ j \in \{cs, hrc, hwg, ct, hx\}, n \in \{1, \dots, N_j\}, t \in \bar{\mathcal{T}} \tag{7.3.9c}$$

$$P_{j,t} = \sum_{n=1}^{N_j} p_{n,j,t},\ j \in \{cs, hrc, hwg, ct, hx\}, t \in \bar{\mathcal{T}} \tag{7.3.9d}$$

$$r_t^e = \sum_{j \in \{cs,hrc,hwg,ct\}} \alpha_j^e P_{j,t} + L_t^e,\ t \in \bar{\mathcal{T}} \tag{7.3.9e}$$

$$r_t^j = \alpha_{i_j}^j P_{i_j,t},\ j \in \{w, ng\},\ i_w = ct, i_{ng} = hwg,\ t \in \bar{\mathcal{T}} \tag{7.3.9f}$$

$$P_{ct,t} = \alpha_{cs}^{cond} P_{cs,t} + P_{hx,t},\ t \in \bar{\mathcal{T}} \tag{7.3.9g}$$

$$P_{cs,t} + P_{hrc,t} + P_{cw,t} + U_{cw,t} = L_t^{cw},\ t \in \bar{\mathcal{T}} \tag{7.3.9h}$$

$$\alpha_{hrc}^h P_{hrc,t} + P_{hwg,t} - P_{hx,t} + P_{hw,t} + U_{hw,t} = L_t^{hw},\ t \in \bar{\mathcal{T}} \tag{7.3.9i}$$

$$E_{j,t+1} = E_{j,t} - P_{j,t},\ j \in \{cw, hw\}, t \in \bar{\mathcal{T}} \tag{7.3.9j}$$

$$0 \le E_{j,t} \le \overline{E}_{j,t},\ j \in \{cw, hw\}, t \in \mathcal{T} \tag{7.3.9k}$$

$$\underline{P}_j \le P_{j,t} \le \overline{P}_j,\ j \in \{cw, hw\}, t \in \bar{\mathcal{T}} \tag{7.3.9l}$$

$$U_{j,t} \ge 0,\ j \in \{cw, hw\}, t \in \bar{\mathcal{T}} \tag{7.3.9m}$$

Here, the binary variables $Y_{n,j,t}$ in (7.3.9b) are used to introduce the logical ON/OFF decisions for each of the chiller, HR chiller, hot water generator, cooling tower, and dump HX units, and constraints (7.3.9c) provide the operating bounds of each of these units

when ON. Constraints (7.3.9d) define the total operating loads of the chiller subplant, HR chiller subplant, hot water generators, cooling towers, and dump HX, respectively. The constraints (7.3.9e)-(7.3.9f) compute the demands of electricity, water, and natural gas (r_t^e, r_t^w, r_t^{ng}) that need to be purchased from the utility companies. Constraints (7.3.9g) impose the energy balance for the condenser water. Constraints (7.3.9h) and (7.3.9i) ensure that the chilled and hot water loads are met. Slack variables $U_{j,k}$, $j \in \{cw, hw\}$ for unmet loads for chilled and hot water ensure that the problem isn't infeasible in case of unmet (under-production) chilled water or hot water. The slack variables for unmet loads are penalized in the objective function by with the penalty ρ_j, $j \in \{cw, hw\}$. Dynamics of the state of charge (SOC) for chilled and hot water storage tanks are given by constraints (7.3.9j). Constraints (7.3.9k)-(7.3.9m) provide bounds on the states, discharge rates from storage tanks, and slack variables for unmet loads. The lower bounds for the discharge rates of chilled water and hot water storage units correspond to the maximum charging rates, which are negative of the maximum discharging rates (i.e, $\underline{P}_j = -\overline{P}_j, j \in \{cw, hw\}$).

Figures 7.2-7.5 show the historical data for the campus electrical load, hot water load, chilled water load, and the electricity prices for the entire year. The vertical red lines in these figures represent monthly periods. In this work, we use these data as the available forecasts of the disturbances in the MPC formulation.

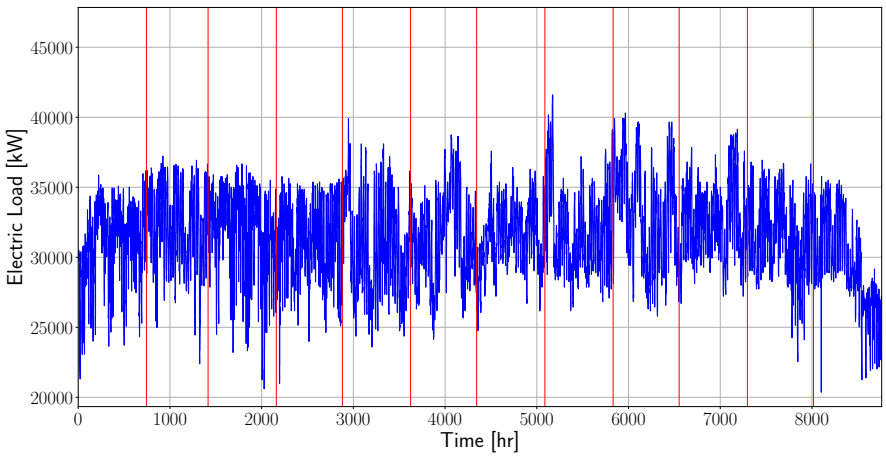

Figure 7.2: Historical electrical load of the campus. Red vertical lines denote end of each month.

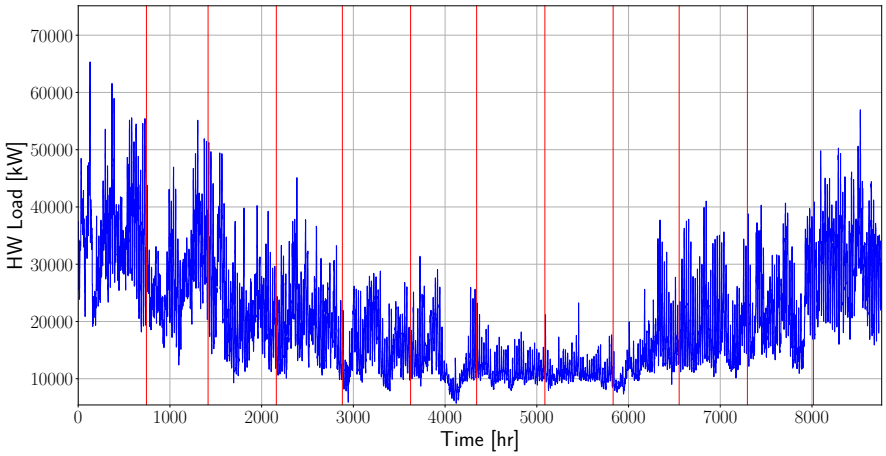

Figure 7.3: Historical hot water load of the campus. Red vertical lines denote end of each month.

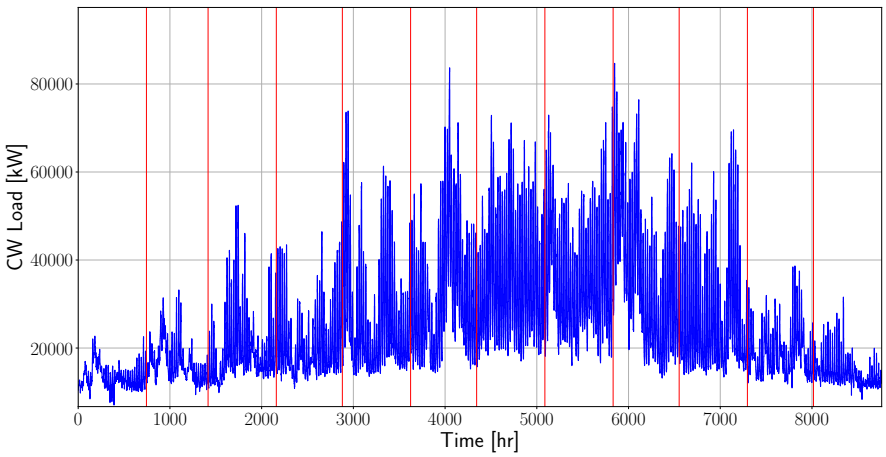

Figure 7.4: Historical chilled water load of the campus. Red vertical lines denote end of each month.

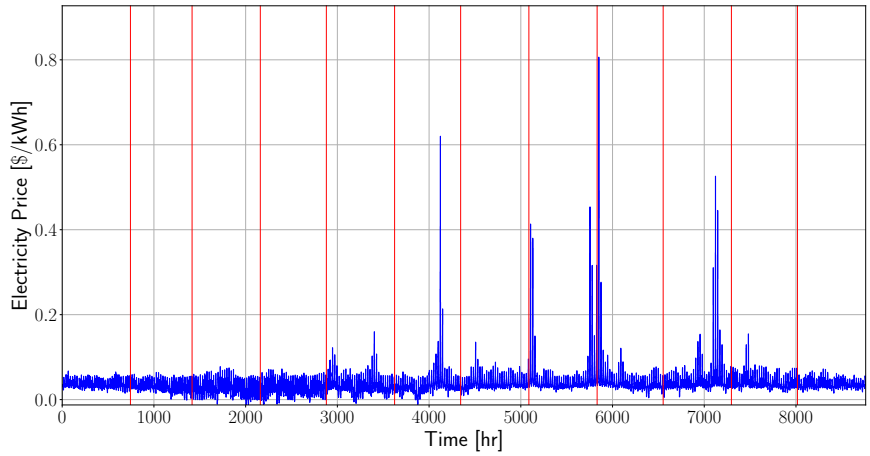

Figure 7.5: Historical electricity price data. Red vertical lines denote end of each month.

7.3.3 *Results*

We now present the results of the computational experiments for the central HVAC plant problem described in Section 7.3.2. First, we implement the dual dynamic programming scheme for the LP relaxation (continuous form) of the MPC problem for the central HVAC

plant, and show that the DDP scheme is able to obtain the optimal solution within an optimality gap of 0.1%. Then, we present the results for the implementation of the DDIP scheme to the mixed-integer MPC for the central HVAC plant with different horizon lengths and discretizations to show that the DDIP scheme provides a scalable approach to solve large-scale and multiscale mixed-integer MPC problems. The schemes are implemented in `Julia` and leverage the algebraic modeling capabilities of `JuMP` (Dunning et al., 2017). The stage-wise optimization problems in the forward and backward sweeps are solved using `Gurobi 9.0`. The algorithms were run on a 32-core machine with Ubuntu 14.04, Intel Xeon 2.30 GHz processors, and 188 GB RAM. The stopping criteria for the DDIP scheme is set to an optimality gap (gap between upper and lower bounds) of 0.5%. For comparisons with the optimal solution for the state evolution, the full problems in the extensive form are also solved.

Continuous (LP Relaxation) MPC Problem

We solve the continuous form or LP relaxation of the MPC problem (7.3.9) with a 168-hour (1-week) horizon. The full problem consists of 9,412 variables and 9,243 linear constraints. Figures 7.6-7.7 present the results for the continuous form (LP relaxation) of the mixed-integer MPC problem for the central HVAC plant. The top and bottom figures in Figure 7.6 show that the DDIP scheme obtains the optimal solution (within an optimality gap of 0.06%) in less than 200 iterations. It takes CPU clock time of 21s to achieve the stopping criteria. We note that the extensive form of the problem is computationally not expensive, and can be solved quite easily.

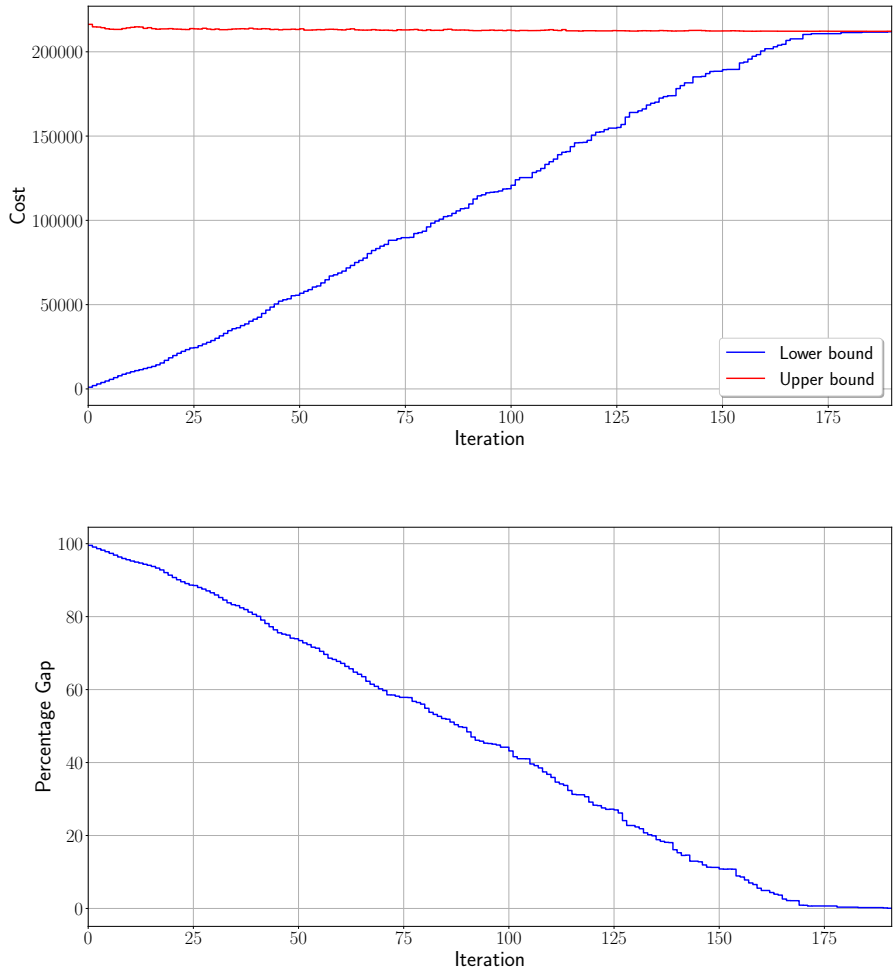

Figure 7.6: **Top:** Evolution of upper and lower bounds over iterations for 168-hour horizon, continuous MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 168-hour horizon, continuous MPC.

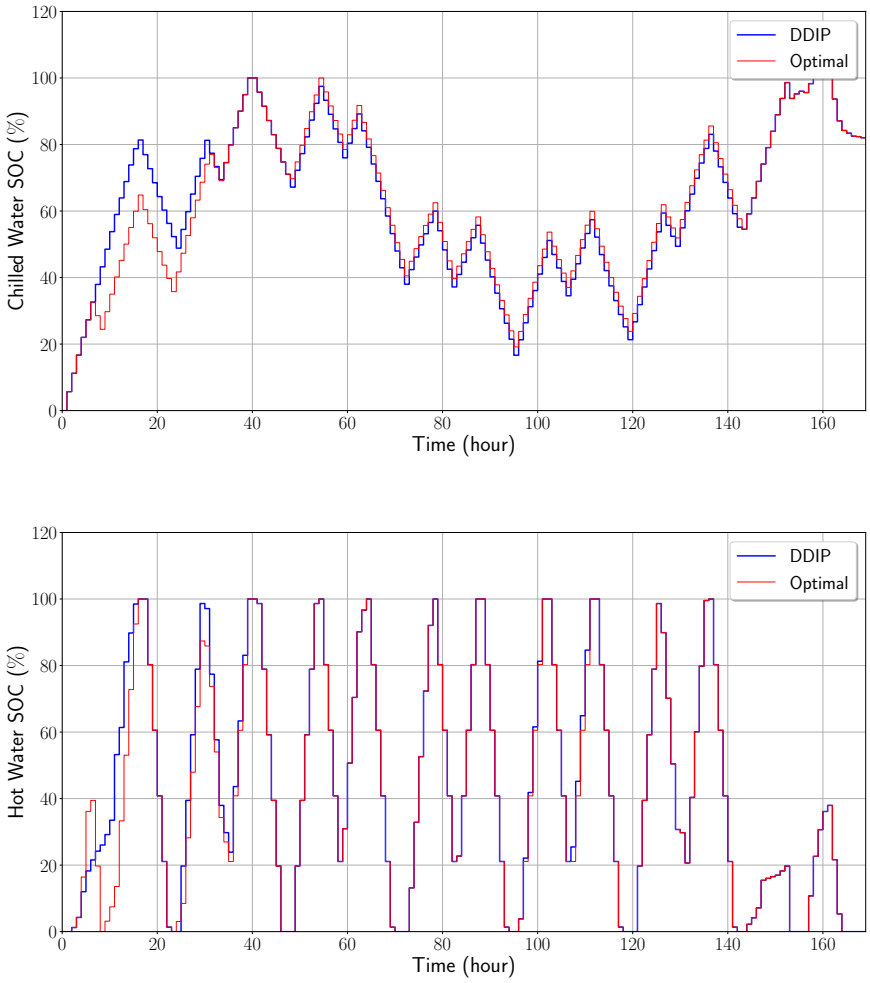

Figure 7.7: Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 168-hour horizon, continuous MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution.

Figure 7.7 provides a comparison of the optimal solution and the final solution from the DDIP scheme for the state of charge evolution of the chilled water and hot water tanks over time. It can be observed that the DDIP scheme obtains close to optimal evolution of the states of charge with slight differences only at a few times which can be attributed to the optimality gap of 0.06%.

Short-Horizon Mixed-integer MPC Problem

We solve the mixed-integer MPC problem for the central HVAC plant with a 168-hour horizon. The full problem consists of 9,412 variables including 3,360 binary variables, and 9,243 linear constraints. Figures 7.8-7.9 present the results for this short-horizon mixed-

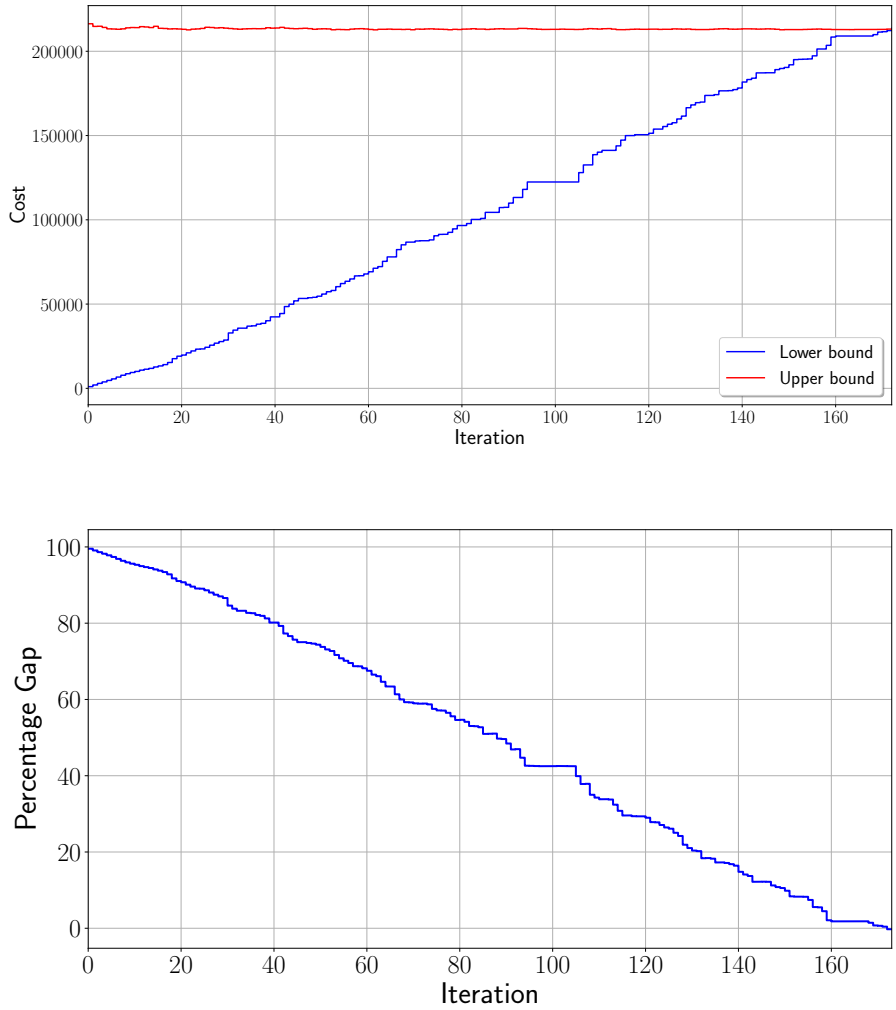

Figure 7.8: **Top:** Evolution of upper and lower bounds over iterations for 168-hour horizon, mixed-integer MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 168-hour horizon, mixed-integer MPC.

integer MPC problem for the central HVAC plant. The top and bottom figures in Figure

7.8 show that the DDIP scheme obtains the optimal solution (within an optimality gap of 0.25%) in approximately 170 iterations. It takes CPU clock time of 56s to achieve the stopping criteria of less than 0.5% optimality gap.

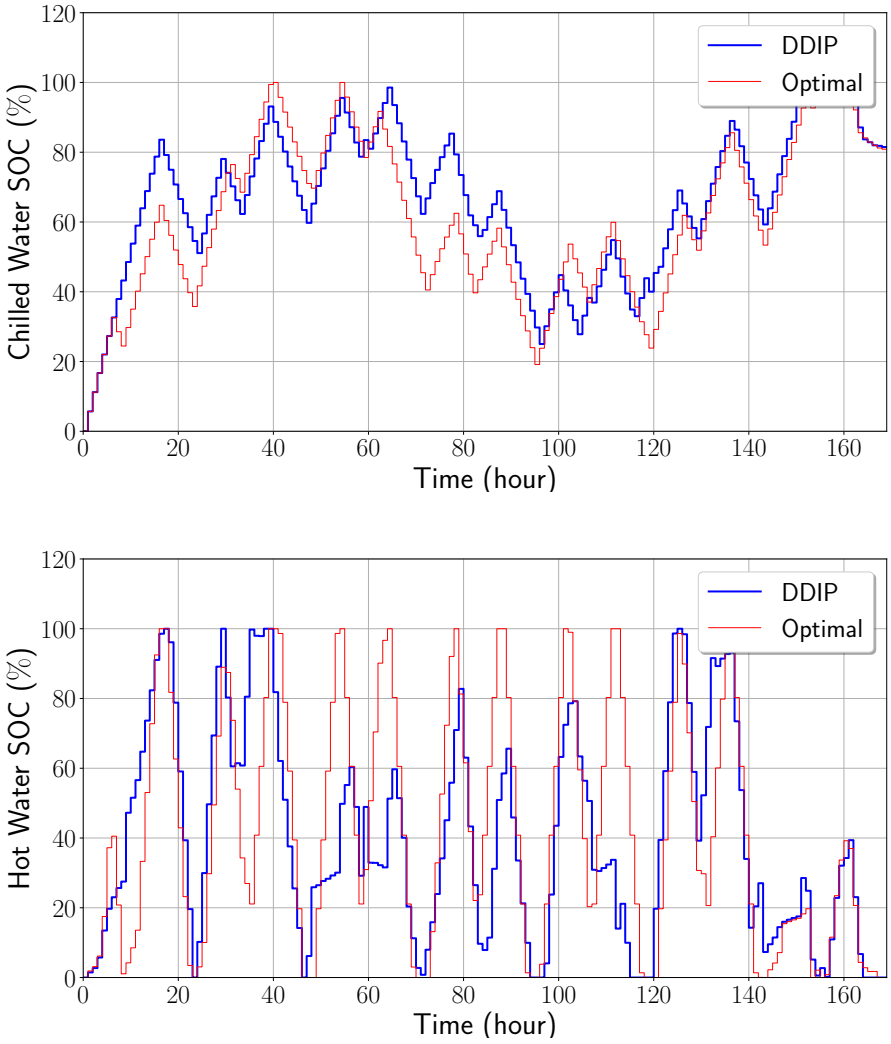

Figure 7.9: Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 168-hour horizon, mixed-integer MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution.

Figure 7.9 provides a comparison of the optimal solution and the final solution from the DDIP scheme for the state of charge evolution of the chilled water and hot water tanks

over time. It can be observed that the DDIP scheme obtains similar to optimal evolution of the states of charge with some deviation from the optimal solution which can be attributed to the optimality gap of 0.25%.

Long-Horizon Mixed-integer MPC

To show that the DDIP scheme is scalable to long-horizon problems, we solve the mixed-integer MPC problem for the central HVAC plant with horizon of 1 month or 720 hours which is more than four times longer horizon than the previous case. The full problem contains 40,324 variables including 1,440 binary variables, and 39,603 linear constraints. The top and bottom figures in Figure 7.10-7.11 present the results for the 720-hour horizon mixed-integer MPC problem for the central HVAC plant. From the top and bottom figures in Figure 7.10, we see that the DDIP scheme achieves an optimality gap of within the stopping criteria (actually achieves a gap of 0.03%) at the end of approximately 710 iterations. Thus, we see that the number of iterations scales approximately with the horizon length. It takes CPU clock time of 256s to achieve the stopping criteria of less than 0.5% optimality gap, which also scales with the horizon length. Since the DDIP scheme needs to solve only single-stage problems in every iteration, the algorithm will easily scale to longer horizons. Also, as noted earlier, a parallel implementation of the backward sweep can further accelerate the computational time for the algorithm for larger problems. Figure 7.11 compares the optimal solution and the final solution from the DDIP scheme for the state of charge evolution of the chilled water and hot water tanks over time. It can be observed that the evolution of the states of charge obtained from DDIP scheme has some deviation from the optimal evolution during some periods while closely following the optimal evolution at other periods. This is due to the fact that the objective of the MPC problem was to minimize the total energy cost and it is possible for different state evolutions to achieve close to optimal objective value. Since the final optimality gap is 0.03%, the state evolution is not the same as optimal, however, the goal of minimizing the energy costs over the time horizon can be achieved. We also highlight that this solution is just the solution of the

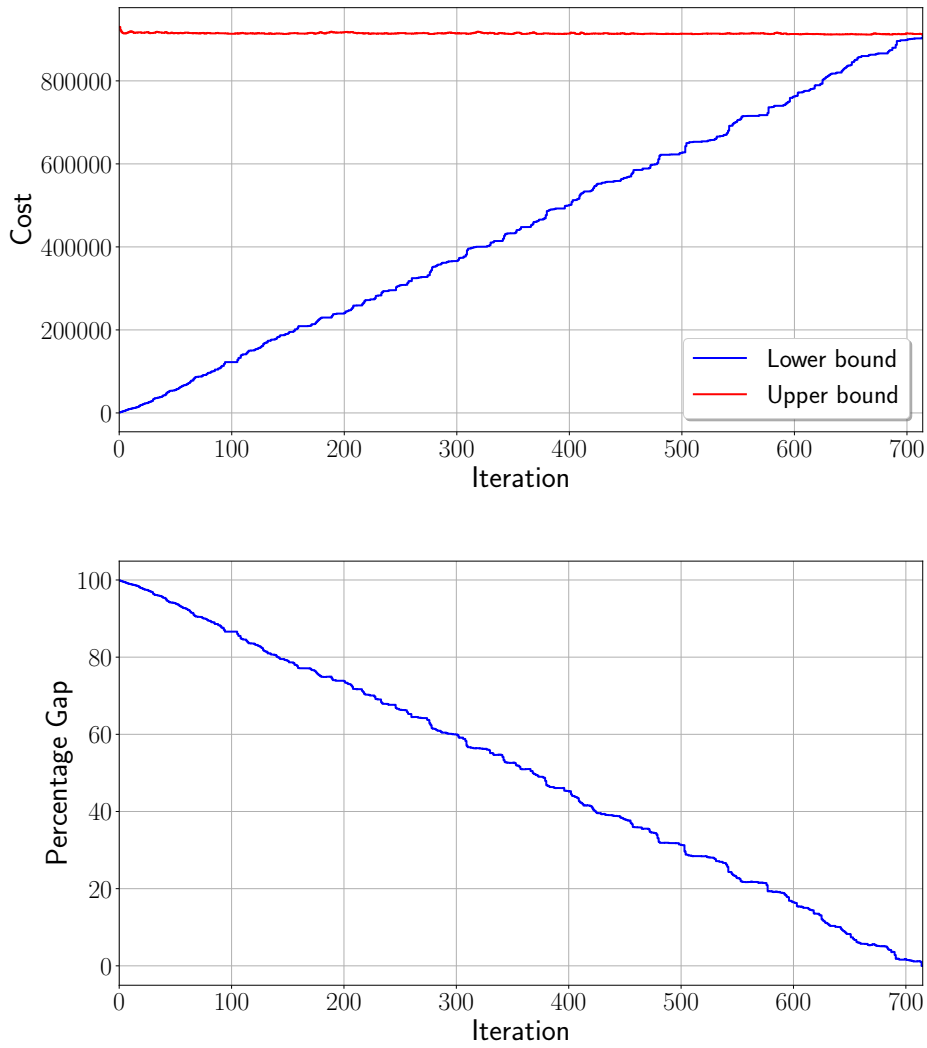

Figure 7.10: **Top:** Evolution of upper and lower bounds over iterations for 720-hour horizon, mixed-integer MPC; red curve represents the evolution of the upper bound and blue curve represents the evolution of the lower bound. **Bottom:** Evolution of the percentage optimality gap over iterations for 720-hour horizon, mixed-integer MPC.

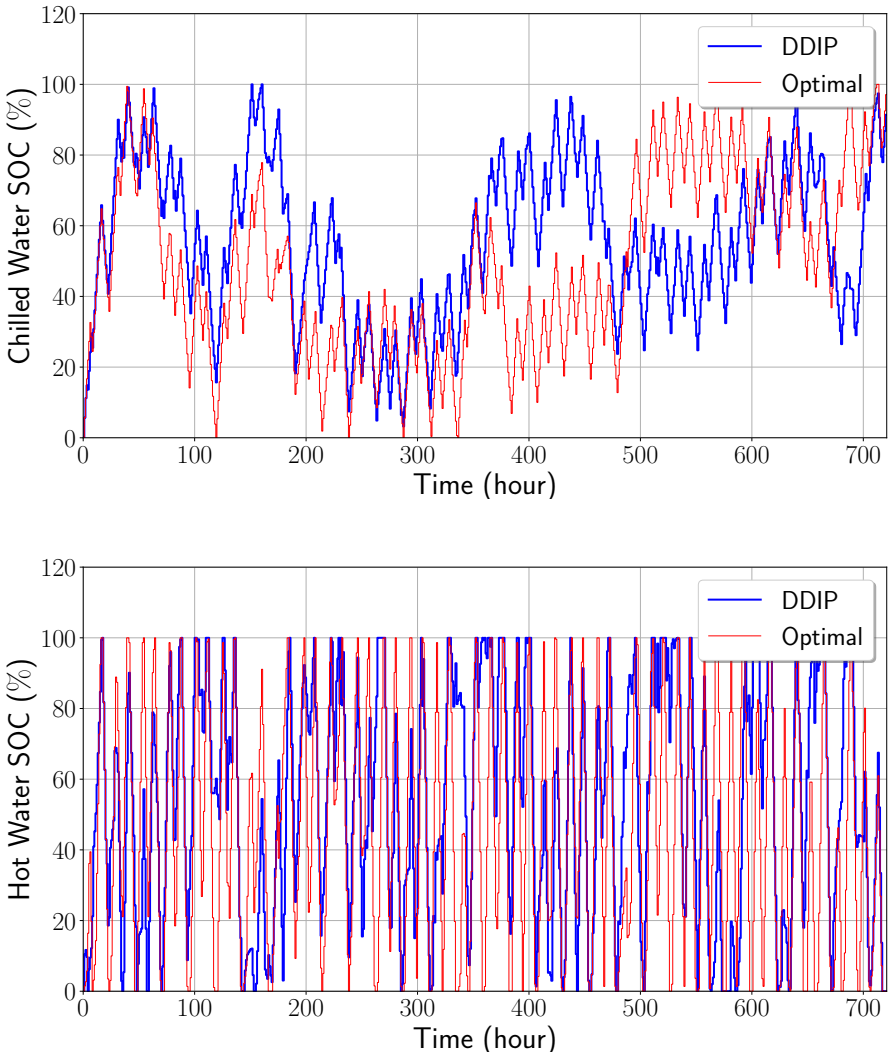

Figure 7.11: Comparison of DDIP and optimal solution of the state of charge evolution for chilled water tank (top) and hot water tank (bottom) over time for 720-hour horizon, mixed-integer MPC. Blue curve represents the solution from the DDIP scheme and red curve represents the optimal solution.

single MPC problem, and the states at the first time step obtained from DDIP algorithm are the same as optimal within a reasonable tolerance. If this scheme is implemented in a receding horizon framework, we can achieve close to optimal state evolution for the entire horizon.

Thus, we conclude that the DDIP scheme is scalable and solves long-horizon large-scale mixed-integer MPC problems to the desired optimality gap. The DDIP scheme is also able to obtain close to optimal evolution of the states within a reasonable tolerance. We highlight that upper bound achieves close to the optimal solution in the initial few iterations because the forward sweeps obtain feasible trajectories of controls and states which give near-optimal objective values, while the lower bound gradually moves close to the optimal solution. This highlights the fact that the lower bound is limited by the quality of the cutting planes obtained in each iteration. We highlight that this characteristic of the lower bound is specific to the problem structure.

7.4 Conclusions and Future Work

We proposed a dual dynamic integer programming (DDIP) scheme for solving multiscale mixed-integer model predictive control (MPC) problems that can handle real constraints of energy systems such as discrete (ON/OFF) controls and minimum startup loads. TO design the DDIP scheme, we leverage the nested cutting plane scheme used in the stochastic dual dynamic programming for multi-stage stochastic programming problems, which uses forward and backward sweeps along the time horizon to adaptively construct and refine the cost-to-go functions. The cutting planes in the DDIP scheme are obtained as the Benders cuts from the LP relaxations of the mixed-integer problems in the backward sweeps. The framework is general and can also be implemented to other applications such online scheduling and integrated production scheduling and control.

We performed computational experiments by implementing the DDIP scheme on a mixed-integer MPC problem for a central HVAC plant for a typical university campus.

We showed that for the continuous MPC problem, the DDIP scheme obtains the optimal solution within an optimality gap of 0.1%, while the implementations of the DDIP scheme to the mixed-integer MPC problem with different horizon lengths and discretizations show that the DDIP scheme provides a scalable approach to solve such large-scale problems.

Future work will focus on implementation of the proposed scheme for stochastic mixed-intger MPC applications with uncertainties evolving over long time horizons and/or fine time resolutions. Parallel implementation of the DDIP scheme to accelerate the algorithm for extremely long horizons or fine time resolutions, and specifically for stochastic mixed-integer problems, is another area of future work.

8

CONCLUSIONS AND FUTURE DIRECTIONS

The major goals of this research were to address the challenges associated with handling uncertainties encountered in the electricity markets, energy demands for the electricity market participation of energy systems. We addressed these challenges by developing a multiscale stochastic model predictive control architecture for energy systems that handles interactions and uncertainties at multiple timescales, and develop an extensive computational framework for the benchmarking of stochastic and deterministic MPC. We also address the challenges of tractability of the large-scale problems arising from long horizons and fine time resolutions by designing stochastic programming based hierarchical MPC and stochastic dual dynamic programming for MPC problems. We also presented implementations of all the computational frameworks developed in this research to industrial case studies. The major contributions of this thesis are summarized in Section 8.1 and some future research directions are discussed in Section 8.2.

8.1 Contributions

8.1.1 *Stochastic MPC Framework for Electricity Market Participation and Demand Charge Mitigation*

In Chapter 2, we proposed a stochastic model predictive control (MPC) framework to determine the optimal battery participation commitments in the energy and frequency regulation markets operated by the ISO and simultaneously mitigate demand charges for an attached modulated load such as buildings. Our stochastic MPC framework solves a two-stage stochastic program that maximizes the expected revenue over a receding horizon and that accounts for the uncertainty in the load, energy and regulation prices, and regulation signals. Our framework also captures the periodicity in the data such as the load and electricity prices by using a Ledoit-Wolf covariance estimator to generate load and price scenario profiles from limited historical data. We benchmark the performance of stochastic MPC against perfect information MPC and deterministic MPC for different prediction horizon lengths and demand charge discounting strategies by using real load data for a typical university campus and price and FR signal data from PJM.

8.1.2 *Computational Framework for Benchmarking Stochastic and Deterministic MPC*

In Chapters Chapter 3 and Chapter 4, we presented an extensive computational framework that integrates forecasting, uncertainty quantification, and model predictive control (MPC) to benchmark the performance of deterministic and stochastic MPC. The computational framework for benchmarking of stochastic and deterministic MPC also implements the detection, quantification, and mitigation of constraint violations that the deterministic MPC can cause due to not accounting for the uncertainty in the data. By means of two case studies, namely a battery management case study and a central HVAC plant for a typical

university campus, we showed that the off-the-shelf deterministic MPC implementations can suffer significant losses in performance and constraint violations due to their inability to handle disturbances that cannot be adequately represented by mean (most likely) forecasts.

8.1.3 *A Retroactive Hierarchical MPC Scheme for Periodic Systems to Handle Long Horizons*

In Chapter 5, we designed a hierarchical model predictive control (MPC) framework using stochastic programming to handle long horizon (or infinite horizon) problems for periodic systems. We showed that an infinite-horizon problem for a periodic system can be cast as a stochastic program (SP) solving a retroactive optimization problem that progressively accumulates historical data and delivers optimal state targets in the long run. The proposed hierarchical MPC scheme consists of the retroactive problem at the high-level hierarchical layer that provides targets to guide a low-level MPC controller that operates over a short period at high time resolution. We also derived a retroactive scheme tailored to linear systems by using cutting plane techniques, analyze stability properties, and suggest strategies to handle nonlinear systems.

8.1.4 *Dual Dynamic Programming Based Schemes for Long-Horizon MPC*

In Chapter 6, we derived and interpreted the stochastic dual dynamic programming (SDDP) for MPC to provide a scalable approach to handle complex MPC applications with uncertainties evolving over long time horizons and with fine time resolutions. The nested cutting-plane scheme of SDDP, which uses forward and backward sweeps along the time horizon to adaptively construct and refine the cost-to-go functions is also leveraged in Chapter 7 to develop a dual dynamic programming based approach for solving multi-scale mixed-integer model predictive control (MPC) problems that can handle real constraints of energy systems such as discrete actuators and minimum startup loads.

8.2 Future research directions

This dissertation gives rise to many future directions in the areas of multiscale control under uncertainty, modeling and control of energy systems, and hierarchical MPC schemes. A direct extension of this research is the parallel implementations of the stochastic MPC and hierarchical MPC schemes presented in this thesis. More future directions arising from this research are discussed below:

8.2.1 *Extensions of the Stochastic MPC Framework to More Complex Market Conditions*

The stochastic MPC framework for energy storage systems in electricity markets for frequency regulation and demand charge mitigation presented in Chapter 2 can be extended to include more complex market conditions. The incorporation of uncertainty models that capture correlations between market prices and loads, and that capture longer timescales is a possible extension of the work. In this case, larger amounts of historical data will be needed to construct the probability distributions of the disturbances and generate realizations. Another direction is to extend the stochastic MPC framework to a multiscale setting in which energy and frequency regulation transactions take place at higher frequencies (minutes to seconds) (Dowling et al., 2017). Some results of stability analysis for receding horizon stochastic MPC are presented in Paulson et al. (2015); Lorenzen et al. (2016). However, for a general stochastic MPC framework with peak costs in the objective function, more analysis is required to provide asymptotic stability and optimality guarantees for an infinite horizon problem.

8.2.2 *Future Directions on Benchmarking Framework for Stochastic MPC*

As part of future work for the computational framework for benchmarking stochastic and deterministic MPC presented in Chapter 3 and Chapter 4, more detailed physics-based

models for battery and units in the central HVAC plants that account for degradation effects explicitly can be explored. Moreover, FR signals at higher time resolutions (seconds) (for the battery management case study) can be included for more detailed studies. Parallel decomposition and simulations will be required in order to accelerate simulations for the resulting large-scale and complex models. The performance obtained with quasi-stochastic MPC also suggests that significant gains in performance can be obtained by using simple uncertainty characterizations, and this approach can be explored in a more general setting. Other important directions of future work involve systematic comparisons of more formulations of model predictive control under uncertainty such as affine decision rules (Domahidi et al., 2012; Guigues and Henrion, 2017), multi-stage and two-stage formulations (Mesbah, 2016), chance-constrained formulations (Mesbah et al., 2014; Farina et al., 2016; Paulson et al., 2017). Recent advances in techniques such as stochastic dual dynamic programming enable the solution of multi-stage problems with many stages (Kumar et al., 2018c). This can enable a systematic comparison of long horizons and number of scenarios on the economic impacts. Another interesting area of research includes scenario reduction and clustering techniques (Römisch, 2009; Kaut and Wallace, 2003; Dupačová et al., 2000; Growe-Kuska et al., 2003) that can be used to enforce feasibility with a small number of samples.

8.2.3 *Extensions of Retroactive Hierarchical MPC Schemes*

An interesting area of research arising from Chapter 5 will be to extend the retroactive hierarchical MPC schemes for periodic systems based on stochastic programming for nonlinear systems and to provide optimality and stability results. In particular, the stochastic programming representation makes it possible to use other schemes such as proximal point methods for convex nonlinear problems Bertsekas (2011). Other future directions are to investigate convergent schemes that have faster convergence than cutting planes, that prevent the accumulation of large amounts of data over time, and that factor in forecast

information in a more effective manner. Another future work in this area is the design of retroactive hierarchical schemes for a more general setting with more than two hierarchical layers. This case will require the implementation of the cutting-plane scheme in a nested fashion to update the state targets at multiple hierarchical layers. The optimality and stability results will also need to be established for this case.

8.2.4 *More Developments in Dual Dynamic Programming Schemes*

Future work for the stochastic dual dynamic programming for multiscale MPC presented in Chapter 6 includes establishing the convergence of the modified SDDP scheme for the multiscale problems. A future investigation is also required to study how to obtain bounds on constraint violations. Parallel implementation of SDDP schemes is also possible to accelerate the convergence of this scheme. In particular, one can explore multiple paths at once to obtain more cutting plane information and accelerate the convergence. The dual dynamic programming for mixed-integer MPC presented in Chapter 7 can be extended to develop schemes for stochastic mixed-integer MPC problems. Convergence analysis for stochastic dual dynamic integer programming is provided in Zou et al. (2016, 2019) for problems that have only binary states. However, convergence and optimality need to be established for the stochastic dual dynamic integer programming scheme for mixed-integer MPC where both control decisions and states can be continuous and integer or binary. Parallel implementations of the DDIP scheme is also an area of future work.

A

NOMENCLATURE

A.1 Nomenclature for Chapter 3

Sets and Indices:

- $\bar{\Xi}$: Set of scenarios in stochastic MPC formulation.
- $\tilde{\Xi}$: Set of scenarios used for validation.
- $\mathcal{T} := \{t+1, t+2, \ldots, t+N\}$: Prediction horizon set, where N is the length.
- $\mathcal{M} := \{1, 2, \ldots, M\}$: Simulation horizon set, where M is the horizon length.
- $\mathcal{H} := \{t-H, t-H+1, \ldots, t\}$: Historical horizon set, with H being the length.
- t: Time instant index.
- ξ: Realization index.

Parameters and Data:

- L_t and $L_t(\xi) \in \mathbb{R}_+$: Attached load and realization of the attached load [kW].
- $\hat{L}_t \in \mathbb{R}_+$: Forecast load [kW].

- π_t^e and $\pi_t^e(\xi) \in \mathbb{R}$: Electricity price and realization [\$/kWh].
- $\hat{\pi}_t^e \in \mathbb{R}$: Forecast electricity price [\$/kWh].
- π_t^f and $\pi_t^f(\xi) \in \mathbb{R}_+$: FR capacity price and realization [\$/kW].
- $\hat{\pi}_t^f \in \mathbb{R}_+$: Forecast FR capacity price [\$/kW].
- α_t and $\alpha_t(\xi) \in [-1, 1]$: FR signal and realization [-].
- $\hat{\alpha}_t \in [-1, 1]$: Forecast FR signal [-].
- $\pi^D \in \mathbb{R}_+$: Demand charge rate [\$/kW].
- $\overline{E} \in \mathbb{R}_+$: Battery capacity [kWh].
- $\overline{P} \in \mathbb{R}_+$: Maximum discharging rate [kW].
- $\underline{P} \in \mathbb{R}_+$: Maximum charging rate [kW].
- $\rho \in \mathbb{R}_+$: FR buffer [kWh/kW].
- $\overline{\Delta P} \in \mathbb{R}_+$: Maximum ramping limit [kW/h].
- $d_{\mathcal{H}}$: Disturbance history over horizon $\mathcal{H}$.

MPC Variables:

- $P_t \in \mathbb{R}$: Net battery charge/discharge rate [kW] .
- $F_t \in \mathbb{R}_+$: FR capacity provided to ISO [kW].
- $r_t \in \mathbb{R}_+$: Net residual load requested from utility [kW].
- $E_t \in \mathbb{R}_+$: Battery state of charge (SOC) [kWh].
- $R_t \in \mathbb{R}_+$: Peak residual load observed up to time t [kW].
- $\Phi^{nobat}(\xi) := \pi^D \max_{t \in \mathcal{M}} L_t(\xi)$: Total cost when battery is not attached [\$].

- $\Phi^{sto}(\xi)$: Total cost for stochastic MPC under the realization ξ [$].
- $\Phi^{det}(\xi)$: Total cost for deterministic MPC under the realization ξ [$].
- $\Phi^{perf}(\xi)$: Total cost for perfect information MPC under the realization ξ [$].
- $\text{VB}^{perf}(\xi)$: Ideal value of battery under perfect information MPC under realization ξ [$].
- $\text{VB}^{sto}(\xi)$: Value of battery under stochastic MPC and realization ξ [$].
- $\text{VB}^{det}(\xi)$: Value of battery under deterministic MPC and realization ξ [$].
- $\text{VSMPC}(\xi)$: Value of stochastic MPC under the realization ξ [$].

A.2 Nomenclature for Chapter 4

Sets and indices:

- $\mathcal{T} := \{t, t+1, t+2, \ldots, t+N-1\}$: Prediction horizon time set, where N is the prediction horizon length, t is the current time.
- $\mathcal{Y} := \{1, 2, \ldots, Y\}$: Planning horizon time set, where Y is the planning (or simulation) horizon length.
- $\mathcal{H} := \{t-H, t-H+1, \ldots, t\}$: Historical horizon set, with H being the length.
- $\mathcal{T}_M := \{t_1, t_2, \ldots, t_{12}\}$: Set of ending time indices for all months.
- t: Time instant index.
- t_m: Ending time index (hour) of month m.
- $\bar{\Xi}$: Set of scenarios in stochastic MPC formulation.
- $\tilde{\Xi}$: Set of scenarios used for validation.

- ξ: Realization index.

Model Parameters and Data:

- π_t^e and $\pi_t^e(\xi) \in \mathbb{R}_+$: Electricity price [\$/kWh] over the time interval $[t, (t+1)]$.
- $\hat{\pi}_t^e \in \mathbb{R}_+$: Forecast electricity price [\$/kWh].
- $\pi_t^w \in \mathbb{R}_+$: Price of water [\$/gal] over the time interval $[t, (t+1)]$.
- $\pi_t^{ng} \in \mathbb{R}_+$: Price of natural gas [\$/kWh] over the time interval $[t, (t+1)]$.
- $\pi^D \in \mathbb{R}_+$: Rate of demand charge [\$/kW].
- $\alpha_{cs}^e \in \mathbb{R}_+$: kW of electricity used by chiller subplant per kW chilled water produced [-].
- $\alpha_{hrc}^e \in \mathbb{R}_+$: kW of electricity used by HR chiller subplant per kW chilled water produced [-].
- $\alpha_{hwg}^e \in \mathbb{R}_+$: kW of electricity used by hot water generator per kW hot water produced [-].
- $\alpha_{ct}^e \in \mathbb{R}_+$: kW of electricity used by cooling towers per kW condenser water input [-].
- $\alpha_{ct}^w \in \mathbb{R}_+$: Gallons of water used by cooling towers per kW condenser water input [-].
- $\alpha_{hwg}^{ng} \in \mathbb{R}_+$: kW of natural gas used by hot water generator per kW hot water produced [-].
- $\alpha_{cs}^{cond} \in \mathbb{R}_+$: kW of condenser water produced by chiller subplant per kW chilled water produced [-].
- $\alpha_{hrc}^h \in \mathbb{R}_+$: kW of hot water produced by HR chiller subplant per kW chilled water produced [-].

- $\rho^{cw} \in \mathbb{R}_+$: Penalty for unmet chilled water load [\$/kWh].
- $\rho^{hw} \in \mathbb{R}_+$: Penalty for unmet hot water load [\$/kWh].
- $\sigma_t := \min\{(M-t)/N, 1\}$: Discounting factor for the monthly demand charge price [-].
- $\beta \in [0,1]$: Storage buffer for the chilled water and hot water TES [-].
- L_t^e and $L_t^e(\xi) \in \mathbb{R}_+$: Electrical load of campus [kW] over the time interval $[t,(t+1)]$.
- $\hat{L}_t^e \in \mathbb{R}_+$: Forecast electrical load [kW].
- L_t^{cw} and $L_t^{cw}(\xi) \in \mathbb{R}_+$: Chilled water load [kW] over the time interval $[t,(t+1)]$.
- $\hat{L}_t^{cw} \in \mathbb{R}_+$: Forecast chilled water load [kW].
- L_t^{hw} and $L_t^{hw}(\xi) \in \mathbb{R}_+$: Hot water load [kW] over the time interval $[t,(t+1)]$.
- $\hat{L}_t^{cw} \in \mathbb{R}_+$: Forecast hot water load [kW].
- $\overline{E}_{cw} \in \mathbb{R}_+$: Energy storage capacity of chilled water energy storage [kWh].
- $\overline{E}_{hw} \in \mathbb{R}_+$: Energy storage capacity of hot water energy storage [kWh].
- $\overline{P}_{cs} \in \mathbb{R}_+$: Maximum load of chiller subplant [kW].
- $\overline{P}_{hrc} \in \mathbb{R}_+$: Maximum load of heat recovery (HR) chiller subplant [kW].
- $\overline{P}_{hwg} \in \mathbb{R}_+$: Maximum load of hot water generator [kW].
- $\overline{P}_{ct} \in \mathbb{R}_+$: Maximum load of cooling towers [kW].
- $\overline{P}_{cw} \in \mathbb{R}_+$: Maximum discharging rate of chilled water energy storage [kW].
- $\overline{P}_{hw} \in \mathbb{R}_+$: Maximum discharging rate of hot water energy storage [kW].
- $R_t \in \mathbb{R}_+$: Peak electrical load observed until time $t \in \mathcal{T}_M$ [kW].

- $v_{cw,t}$: Random variable used to update chilled water storage from predicted value to simulated actual value.
- $v_{hw,t}$: Random variable used to update hot water storage from predicted value to simulated actual value.
- $\sigma^2_{cw,err,t+1}$: Variance of chilled water load prediction error for $t+1$.
- $\sigma^2_{hw,err,t+1}$: Variance of hot water load prediction error for $t+1$.
- $\sigma^2_{cw,int}$: Variance of integrated chilled water load for 1-hour periods.
- $\sigma^2_{hw,int}$: Variance of integrated hot water load for 1-hour periods.

Controls:

- $P_{cs,t} \in \mathbb{R}_+$: Amount of chilled water produced by chiller subplant [kW] over the time interval $[t, (t+1)]$.
- $P_{hrc,t} \in \mathbb{R}_+$: Amount of chilled water produced by heat recovery (HR) chiller subplant [kW] over the time interval $[t, (t+1)]$.
- $P_{hwg,t} \in \mathbb{R}_+$: Amount of hot water produced by hot water generator [kW] over the time interval $[t, (t+1)]$.
- $P_{ct,t} \in \mathbb{R}_+$: Amount of condenser water input to the cooling towers [kW] over the time interval $[t, (t+1)]$.
- $P_{cw,t} \in \mathbb{R}$: Net charge/discharge rate [kW] of the chilled water energy storage over the time interval $[t, (t+1)]$. If $P_{cw,t} > 0$, the chilled water is being discharged and if $P_{cw,t} < 0$ the chilled water is being charged.
- $P_{hw,t} \in \mathbb{R}$: Net charge/discharge rate [kW] of the hot water energy storage over the time interval $[t, (t+1)]$. If $P_{hw,t} > 0$, the hot water is being discharged and if $P_{hw,t} < 0$ the hot water is being charged.

- $P_{hx,t} \in \mathbb{R}_+$: Amount of hot water input to the dump heat exchanger (HX) [kW] over the time interval $[t, (t+1)]$.
- $r_t^e \in \mathbb{R}$: Residual electrical load [kW] over the time interval $[t, (t+1)]$.
- $r_t^w \in \mathbb{R}$: Residual water demand [gal/h] over the time interval $[t, (t+1)]$.
- $r_t^{ng} \in \mathbb{R}$: Residual natural gas demand [kW] over the time interval $[t, (t+1)]$.
- $S_{cw,t}^{un} \in \mathbb{R}$: Slack variable for unmet chilled water load [kW] over the time interval $[t, (t+1)]$.
- $S_{cw,t}^{ov} \in \mathbb{R}$: Slack variable for overmet (over-produced) chilled water load [kW] over the time interval $[t, (t+1)]$.
- $S_{hw,t}^{un} \in \mathbb{R}$: Slack variable for unmet hot water load [kW] over the time interval $[t, (t+1)]$.
- $S_{hw,t}^{ov} \in \mathbb{R}$: Slack variable for overmet (over-produced) hot water load [kW] over the time interval $[t, (t+1)]$.

States:

- $E_{cw,t} \in \mathbb{R}_+$: Energy level of the chilled water energy storage [kWh] at time t.
- $E_{hw,t} \in \mathbb{R}_+$: Energy level of the hot water energy storage [kWh] at time t.
- R_t: Peak residual electrical load observed up to time t [kW].
- $ul_{cw,t} \in \mathbb{R}$: State variable integrating the slack variable for unmet chilled water load [kW] over the time interval $[t, (t+1)]$.
- $ol_{cw,t} \in \mathbb{R}$: State variable integrating the slack variable for overmet (over-produced) chilled water load [kW] over the time interval $[t, (t+1)]$.
- $ul_{hw,t} \in \mathbb{R}$: State variable integrating the slack variable for unmet hot water load [kW] over the time interval $[t, (t+1)]$.

- $ol_{hw,t} \in \mathbb{R}$: State variable integrating the slack variable for overmet (over-produced) hot water load [kW] over the time interval $[t, (t+1)]$.

Economic metrics:

- $\Phi^{nocp}(\xi) := \pi^D \max_{t \in \mathcal{M}} L_t(\xi)$: Total cost when there is no HVAC central plant in campus [$].
- $\Phi^{sto}(\xi)$: Total cost for stochastic MPC under the realization ξ [$].
- $\Phi^{det}(\xi)$: Total cost for deterministic MPC under the realization ξ [$].
- $\Phi^{perf}(\xi)$: Total cost for perfect information MPC under the realization ξ [$].
- $\text{CCP}^{perf}(\xi)$: Ideal cost of central plant under perfect information MPC under realization ξ [$].
- $\text{CCP}^{sto}(\xi)$: Cost of central plant under stochastic MPC and realization ξ [$].
- $\text{CCP}^{det}(\xi)$: Cost of central plant under deterministic MPC and realization ξ [$].
- $\text{VSMPC}(\xi)$: Value of stochastic MPC under the realization ξ [$].

A.3 Nomenclature for Chapter 7

Sets and indices:

- $\mathcal{T} := \{0, 1, \ldots, N\}$: Prediction horizon time set for the state evolution, where N is the prediction horizon length.
- $\bar{\mathcal{T}} := \{0, 1, \ldots, N-1\}$: Prediction horizon time set for the control trajectory, where N is the prediction horizon length.
- t: Time instant index.

Model Parameters and Data:

- $\pi_t^e \in \mathbb{R}_+$: Electricity price [\$/kWh] over the time interval $[t, (t+1)]$.
- $\pi_t^w \in \mathbb{R}_+$: Price of water [\$/gal] over the time interval $[t, (t+1)]$.
- $\pi_t^{ng} \in \mathbb{R}_+$: Price of natural gas [\$/kWh] over the time interval $[t, (t+1)]$.
- $\alpha_{cs}^e \in \mathbb{R}_+$: kW of electricity used by each chiller unit per kW chilled water produced [-].
- $\alpha_{hrc}^e \in \mathbb{R}_+$: kW of electricity used by each HR chiller unit per kW chilled water produced [-].
- $\alpha_{hwg}^e \in \mathbb{R}_+$: kW of electricity used by each hot water generator per kW hot water produced [-].
- $\alpha_{ct}^e \in \mathbb{R}_+$: kW of electricity used by each cooling tower per kW condenser water input [-].
- $\alpha_{ct}^w \in \mathbb{R}_+$: Gallons of water used by each cooling tower per kW condenser water input [-].
- $\alpha_{hwg}^{ng} \in \mathbb{R}_+$: kW of natural gas used by each hot water generator per kW hot water produced [-].
- $\alpha_{cs}^{cond} \in \mathbb{R}_+$: kW of condenser water produced by each chiller unit per kW chilled water produced [-].
- $\alpha_{hrc}^h \in \mathbb{R}_+$: kW of hot water produced by each HR chiller unit per kW chilled water produced [-].
- $\rho^{cw} \in \mathbb{R}_+$: Penalty for unmet chilled water load [\$/kWh].
- $\rho^{hw} \in \mathbb{R}_+$: Penalty for unmet hot water load [\$/kWh].
- $\beta_{cs} \in [0, 1] = 0.5$: Minimum operating load (fraction of maximum capacity) of each chiller unit when ON.

- $\beta_{hrc} \in [0,1] = 0.8$: Minimum operating load (fraction of maximum capacity) of each HR chiller unit when ON.
- $\beta_{hwg} \in [0,1] = 0.5$: Minimum operating load (fraction of maximum capacity) of each hot water generator when ON.
- $\beta_{ct} \in [0,1] = 0.5$: Minimum operating load (fraction of maximum capacity) of each cooling tower when ON.
- $\beta_{cs} \in [0,1] = 0$: Minimum operating load (fraction of maximum capacity) of dump heat exchanger when ON.
- $L_t^e \in \mathbb{R}_+$: Electrical load of campus [kW] over the time interval $[t,(t+1)]$.
- $L_t^{cw} \in \mathbb{R}_+$: Chilled water load [kW] over the time interval $[t,(t+1)]$.
- $L_t^{hw} \in \mathbb{R}_+$: Hot water load [kW] over the time interval $[t,(t+1)]$.
- $\overline{E}_{cw} \in \mathbb{R}_+$: Energy storage capacity of chilled water energy storage [kWh].
- $\overline{E}_{hw} \in \mathbb{R}_+$: Energy storage capacity of hot water energy storage [kWh].
- $\overline{p}_{cs} \in \mathbb{R}_+$: Maximum load of each chiller unit [kW].
- $\overline{p}_{hrc} \in \mathbb{R}_+$: Maximum load of each heat recovery (HR) chiller unit [kW].
- $\overline{p}_{hwg} \in \mathbb{R}_+$: Maximum load of each hot water generator [kW].
- $\overline{p}_{ct} \in \mathbb{R}_+$: Maximum load of each cooling tower [kW].
- $\overline{P}_{cw} \in \mathbb{R}_+$: Maximum discharging rate of chilled water energy storage [kW].
- $\overline{P}_{hw} \in \mathbb{R}_+$: Maximum discharging rate of hot water energy storage [kW].

Controls:

- $Y_{n,cs,t} \in \{0,1\}$: Binary variable for n^{th} chiller unit: ON if 1, OFF if 0 over the time interval $[t,(t+1)]$.

- $Y_{n,hrc,t} \in \{0,1\}$: Binary variable for n^{th} heat recovery (HR) chiller unit: ON if 1, OFF if 0 over the time interval $[t,(t+1)]$.
- $Y_{n,hwg,t} \in \{0,1\}$: Binary variable for n^{th} hot water generator: ON if 1, OFF if 0 over the time interval $[t,(t+1)]$.
- $Y_{n,ct,t} \in \{0,1\}$: Binary variable for n^{th} cooling tower: ON if 1, OFF if 0 over the time interval $[t,(t+1)]$.
- $Y_{n,hx,t} \in \{0,1\}$: Binary variable for n^{th} dump heat exchanger (HX) unit: ON if 1, OFF if 0 over the time interval $[t,(t+1)]$.
- $p_{n,cs,t} \in \mathbb{R}_+$: Amount of chilled water produced by n^{th} chiller unit [kW] over the time interval $[t,(t+1)]$.
- $p_{n,hrc,t} \in \mathbb{R}_+$: Amount of chilled water produced by n^{th} heat recovery (HR) chiller unit [kW] over the time interval $[t,(t+1)]$.
- $p_{n,hwg,t} \in \mathbb{R}_+$: Amount of hot water produced by n^{th} hot water generator [kW] over the time interval $[t,(t+1)]$.
- $p_{n,ct,t} \in \mathbb{R}_+$: Amount of condenser water input to the n^{th} cooling tower [kW] over the time interval $[t,(t+1)]$.
- $p_{n,hx,t} \in \mathbb{R}_+$: Amount of hot water input to the n^{th} dump heat exchanger (HX) [kW] over the time interval $[t,(t+1)]$.
- $P_{cs,t} \in \mathbb{R}_+$: Total amount of chilled water produced by chiller subplant (4 chiller units) [kW] over the time interval $[t,(t+1)]$.
- $P_{hrc,t} \in \mathbb{R}_+$: Total amount of chilled water produced by heat recovery (HR) chiller subplant (3 HR chiller units) [kW] over the time interval $[t,(t+1)]$.
- $P_{hwg,t} \in \mathbb{R}_+$: Total amount of hot water produced by 3 hot water generators [kW] over the time interval $[t,(t+1)]$.

- $P_{hx,t} \in \mathbb{R}_+$: Total amount of hot water input to the dump heat exchanger (HX) [kW] over the time interval $[t, (t+1)]$.
- $P_{cw,t} \in \mathbb{R}$: Net charge/discharge rate [kW] of the chilled water energy storage over the time interval $[t, (t+1)]$. If $P_{cw,t} > 0$, the chilled water is being discharged and if $P_{cw,t} < 0$ the chilled water is being charged.
- $P_{hw,t} \in \mathbb{R}$: Net charge/discharge rate [kW] of the hot water energy storage over the time interval $[t, (t+1)]$. If $P_{hw,t} > 0$, the hot water is being discharged and if $P_{hw,t} < 0$ the hot water is being charged.
- $r_t^e \in \mathbb{R}$: Residual electrical load [kW] over the time interval $[t, (t+1)]$.
- $r_t^w \in \mathbb{R}$: Residual water demand [gal/h] over the time interval $[t, (t+1)]$.
- $r_t^{ng} \in \mathbb{R}$: Residual natural gas demand [kW] over the time interval $[t, (t+1)]$.
- $U_{cw,t} \in \mathbb{R}$: Slack variable for unmet chilled water load [kW] over the time interval $[t, (t+1)]$.
- $U_{hw,t} \in \mathbb{R}$: Slack variable for unmet hot water load [kW] over the time interval $[t, (t+1)]$.

States:

- $E_{cw,t} \in \mathbb{R}_+$: Energy level of the chilled water energy storage [kWh] at time t.
- $E_{hw,t} \in \mathbb{R}_+$: Energy level of the hot water energy storage [kWh] at time t.

B

HANDLING LONG HORIZONS IN MPC: A STOCHASTIC PROGRAMMING APPROACH

This appendix presents a formulation and preliminary case study of a building system with stationary battery storage to show that the hierarchical MPC scheme proposed in Chapter 5 is optimal under nominal (perfect forecast) conditions and can be extended to handle imperfect forecasts by correcting the targets in real-time. We also provide a comparison of the performance of the proposed hierarchical MPC scheme with standard MPC schemes.

B.1 Introduction

MPC provides a powerful planning framework for industrial applications but it is well-known that its scope is hindered by the computational complexity associated with the length of the planning horizon. This limitation is relevant in applications such as energy systems, inventory management, and scheduling. In the context of energy systems, long planning horizons are often needed to capture monthly peak electricity costs (demand charges) (Ma et al., 2012a). Peak electricity demands remain a serious concern for utilities because they pose capacity challenges to the power grid. This situation has created

incentives for energy storage (de Salis et al., 2014; Joshi and Pindoriya, 2015). In this context, MPC has been used in Braun (1990) to mitigate demand charges of a building HVAC system. Due to the inability to handle long horizons, the authors penalize the peak electricity demand over the short-term receding horizon. Unfortunately, this approach results in overly conservative MPC policies and deteriorating economic performance. Long-term discounting cost factors are used in Patel et al. (2016) to reduce the conservatism of short-term MPC, but this approach is ad-hoc and does not have optimality guarantees.

The work in Zavala (2016) proposes a hierarchical MPC scheme to handle long time horizons. Here, adjoint information obtained from a long-term and coarse MPC controller is used to guide the control policy of a short-term MPC controller. It is demonstrated that this approach can achieve optimality. Unfortunately, such an approach requires continuity of the adjoint profiles, which is not guaranteed in general formulations with peak costs (such as those arising from demand charges) and integer decisions (as those arising in scheduling formulations). Alternative hierarchical MPC schemes (Scattolini and Colaneri, 2007; Picasso et al., 2010) available in the literature provide feasibility but not optimality guarantees.

In this work, we show that when periodicity constraints are enforced over short-term stages, the long-horizon MPC problem can be cast as a stochastic programming (SP) problem. Under this setting, a short-term planning stage represents a scenario, the design variables are the initial/terminal periodic states, and the recourse decisions are the intra-stage operational policies. The SP representation suggests a hierarchical MPC scheme in which optimal design decisions are used as targets to guide a short-term MPC controller. We show that, under nominal conditions with perfect forecasts, the hierarchical scheme yields an optimal policy of the long-term MPC problem. The SP setting also reveals strategies to correct targets and short-term policies when forecasts are imperfect. The proposed approach can be easily generalized to handle peak costs, by incorporating an additional design variable that imposes a demand budget to the short-term MPC controller. The approach can also handle general nonlinear models and integer decisions and can enable the

use of parallel decomposition schemes.

We demonstrate the performance of the hierarchical MPC scheme using an application in buildings with electricity storage. Here, the goal is to determine the optimal short-term (hourly) participation strategies in frequency regulation (FR) and energy markets while simultaneously mitigating long-term (monthly) demand charges from utilities. We compare the solution of the hierarchical MPC scheme with that from a standard MPC scheme that uses a discounting (weighting) factor to account for the inability of solving the long horizon problem. We confirm that the proposed hierarchical MPC scheme achieves optimal policies under perfect forecasts of the electrical load. Under imperfect forecasts, we find that the cost from the hierarchical MPC scheme provides superior performance over standard MPC because it can capture long-term variability. The proposed developments can also be used in other applications that need to handle long horizons, such as inventory management and scheduling formulations.

B.2 Hierarchical MPC Scheme

We consider an MPC planning problem of the form:

$$\min_{u_t} \sum_{t\in\mathcal{T}} \varphi_1(x_t, u_t, d_t) + \max_{t\in\mathcal{T}} \varphi_2(x_t, u_t, d_t) \tag{B.2.1a}$$

$$\text{s.t. } x_{t+1} = f(x_t, u_t, d_t),\ t \in \bar{\mathcal{T}} \tag{B.2.1b}$$

$$x_0 = \bar{x}_0,\ x_t \in \mathcal{X}, u_t \in \mathcal{U}. \tag{B.2.1c}$$

Here, $\mathcal{T} := \{0, \ldots, N\}$ is the set of time steps with N being the final time in the horizon (with $\bar{\mathcal{T}} = \mathcal{T} \setminus \{N\}$), $\varphi_1(\cdot)$ is a time-additive cost function, and $\varphi_2(\cdot)$ is a time-max (peak) cost function. The controls, states, and true (realized) disturbances at time t are denoted as u_t, x_t, and d_t, respectively. The forecasted disturbance at time t is $\hat{d}_t$ and initial state is $\bar{x}_0$.

We consider a partition (in lexicographic order) of the time horizon set $\mathcal{T}$ (and of $\bar{\mathcal{T}}$) into M *stages* with shorter horizons. We denote each stage by $\xi \in \Xi$ with $\mathcal{T}_\xi := \{0, \ldots, N_\xi\}$

being the inner stage time horizon and $\Xi := \{0, \dots, M\}$ being the set of all stages, satisfying $\mathcal{T} = \cup_{\xi \in \Xi} \mathcal{T}_\xi$ and $\sum_{\xi \in \Xi} N_\xi = N$. For convenience, we define the set $\bar{\Xi} := \Xi \setminus \{M\}$. Note that each stage Ξ can have a different time horizon. We partition states, controls, and disturbance policies into stages and denote the stage policies as $u_{\xi,t}$, $x_{\xi,t}$, and $d_{\xi,t}$ for $\xi \in \Xi$ and $t \in \mathcal{T}_\xi$. We use these partitions to reformulate the MPC problem in the following *equivalent* form:

$$\min_{u_{\xi,t}} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) + \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \tag{B.2.2a}$$

$$\text{s.t. } x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}),\ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \tag{B.2.2b}$$

$$x_{\xi+1,0} = x_{\xi,N_\xi}, \xi \in \bar{\Xi} \tag{B.2.2c}$$

$$x_{0,0} = \bar{x}_0,\ x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{B.2.2d}$$

Here, the constraint (B.2.2c) enforces continuity between stages. We modify the MPC planning problem by assuming that periodicity is enforced at the end of every stage:

$$\min_{u_{\xi,t}, x_{\xi,0}} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} \varphi_1(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) + \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} \varphi_2(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}) \tag{B.2.3a}$$

$$\text{s.t. } x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}),\ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \tag{B.2.3b}$$

$$x_{\xi+1,0} = x_{\xi,N_\xi}, \xi \in \bar{\Xi} \tag{B.2.3c}$$

$$x_{\xi,N_\xi} = x_{\xi,0}, \xi \in \Xi \tag{B.2.3d}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{B.2.3e}$$

In this formulation (that we call *long-term MPC*), $x_{0,0}$ is a free variable that we seek to optimize. We also note that the periodicity constraint (B.2.3d) together with the stage continuity constraints (B.2.3c) can be expressed as $x_{\xi+1,0} = x_{\xi,0}$, $\xi \in \bar{\Xi}$. These constraints, in turn, can be reformulated as $x_{\xi,0} = x_0$, $\xi \in \Xi$ by introducing an additional *lifting* variable x_0. Consequently, the goal of the long-term MPC formulation (B.2.3) is to find the optimal periodic state x_0 and control policies $u_{\xi,t}$, $\xi \in \Xi, t \in \mathcal{T}_\xi$ that minimize the time-additive and

peak costs. The assumption on periodicity is reasonable in applications such as energy and inventory management systems, in which disturbance profiles (e.g, demands and prices) have a strong periodic component (Huang et al., 2011; Subramanian et al., 2014; Risbeck et al., 2015).

We can reformulate the peak cost function to obtain the following equivalent form of (B.2.3):

$$\min_{u_{\xi,t},x_0,\eta} \sum_{\xi\in\Xi}\sum_{t\in\mathcal{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \eta \tag{B.2.4a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \le \eta, \xi\in\Xi, t\in\mathcal{T}_\xi \tag{B.2.4b}$$

$$x_{\xi,t+1} = f(x_{\xi,t},u_{\xi,t},d_{\xi,t}),\ \xi\in\Xi, t\in\bar{\mathcal{T}}_\xi \tag{B.2.4c}$$

$$x_{\xi,0} = x_0, \xi\in\Xi \tag{B.2.4d}$$

$$x_{\xi,N_\xi} = x_0, \xi\in\Xi \tag{B.2.4e}$$

$$x_{\xi,t}\in\mathcal{X}, u_{\xi,t}\in\mathcal{U}. \tag{B.2.4f}$$

We denote the solution of (B.2.4) as $x^*_{\xi,t}, u^*_{\xi,t}, \eta^*$ and note that $\eta^* = \max_{\xi\in\Xi}\max_{t\in\mathcal{T}_\xi} \varphi_2(x^*_{\xi,t},u^*_{\xi,t},d_{\xi,t})$ holds. Consequently, η^* is the peak cost over the entire planning horizon. We also note that $x^*_{\xi+1,0} = x^*_{\xi,0} = x^*_0$. From the structure of (B.2.4) we note that the only coupling between stages arises from the variables x_0 and η. Consequently, (B.2.3) can be seen as a *stochastic programming problem* in which stages are operational scenarios, x_0 and η are *design* variables, and $x_{\xi,t}, u_{\xi,t}$ are *scenario* (recourse) policies. When no peak costs are present, the only design variable of the long-term MPC formulation is the periodic state x_0.

By fixing the design variables to their optimal values x^*_0 and η^*, we can decompose problem (B.2.4) into M stage subproblems of the form:

$$\min_{u_{\xi,t}} \sum_{t\in\mathcal{T}_\xi} \varphi_1(x_{\xi,t},u_{\xi,t},d_{\xi,t}) + \eta^* \tag{B.2.5a}$$

$$\text{s.t. } \varphi_2(x_{\xi,t},u_{\xi,t},d_{\xi,t}) \le \eta^*,\ t\in\mathcal{T}_\xi \tag{B.2.5b}$$

$$x_{\xi,t+1} = f(x_{\xi,t}, u_{\xi,t}, d_{\xi,t}),\ t \in \bar{\mathcal{T}}_\xi \tag{B.2.5c}$$

$$x_{\xi,0} = x_0^* \tag{B.2.5d}$$

$$x_{\xi,N_\xi} = x_0^* \tag{B.2.5e}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{B.2.5f}$$

The subproblem has the structure of a standard optimal control problem with periodicity constraints. The SP formulation (B.2.3) thus suggests a *hierarchical MPC scheme*, in which the long-term MPC problem (B.2.4) (equivalently (B.2.2)) guides a short-term MPC controller (with formulation (B.2.5)). The communication between the hierarchical levels arises in the form of targets for the periodic state x_0^* and the peak cost η^*. Figure B.1 provides a sketch of the hierarchical MPC scheme. Because the targets $x_{\xi,0}^*$ and η^* are optimal, the solution of the MPC stage problem yields an optimal stage trajectory $u_{\xi,t}^*$ and $x_{\xi,t}^*$ (or a trajectory that achieves the same optimal stage cost). Consequently, the targets obtained from the SP problem (B.2.2) are optimal when the *disturbance forecasts are perfect* ($d_{\xi,t} = \hat{d}_{\xi,t}$) and the solution of the SP problem is equivalent to that of the long-horizon MPC problem (B.2.3).

The SP setting also indicates that, when the forecasts are imperfect (i.e., $\hat{d}_{\xi,t} \neq d_{\xi,t}$), the targets x_0^*, η^* will not be optimal and the stage (recourse) problem (B.2.5) might not have a feasible solution. This situation can be mitigated by penalizing deviations from the target (which will find the closest feasible point) or by re-optimizing the design targets using the actual realized disturbances when they become available in real-time. The SP setting also reveals that by increasing the number of scenarios in the planning problem (e.g., by using multiple possible realizations of the forecast), it is possible to find targets that remain optimal and feasible for all stage realizations. Consequently, the proposed approach provides a framework to easily construct robust formulations. Another important feature of the hierarchical MPC approach is that it avoids the need to weight or discount the long-term costs in the short-horizon MPC subproblem. Moreover, one can solve the SP

problem efficiently by using decomposition schemes based on parallel linear algebra and Benders/Lagrangian decomposition (Zavala et al., 2008; Geoffrion, 1972).

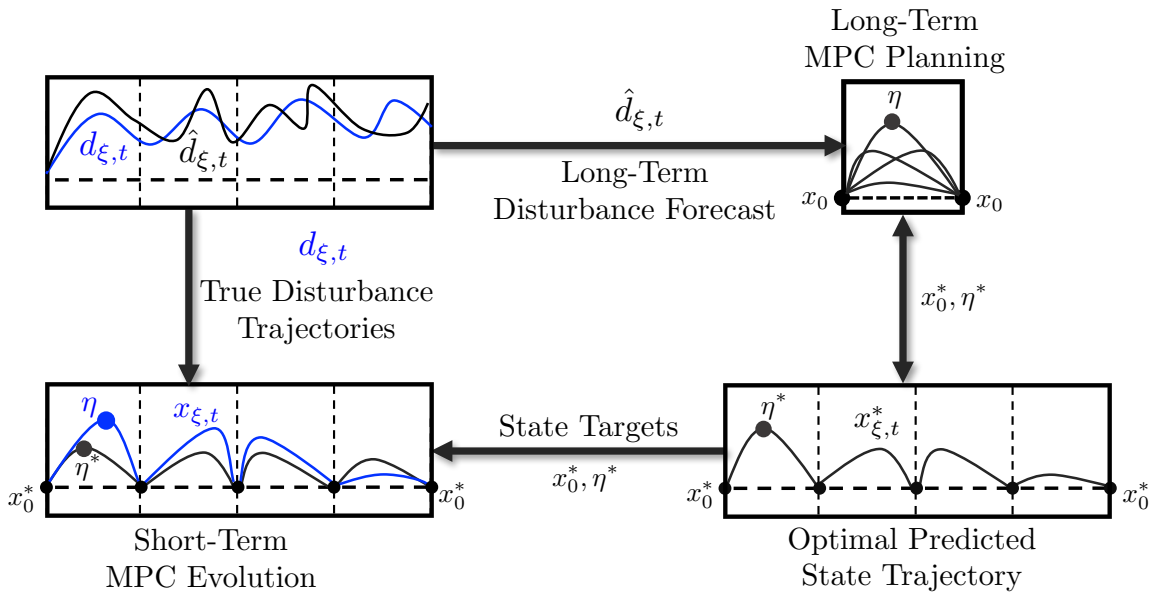

Figure B.1: Hierarchical MPC scheme.

B.3 Stationary Battery Study

Batteries are flexible assets that can be used to provide energy and frequency regulation (FR) capacity for independent system operators (ISOs) and that can aid utility companies by providing demand-side management capabilities for buildings or manufacturing facilities (Rastler, 2010). The use of batteries for *simultaneous* FR and demand charge mitigation has been studied in Shi et al. (2017). In such settings, the objective function of the MPC scheme takes the form of the formulation in (B.2.1), where the additive costs represent market revenue and peak costs represent demand charges. The proposed battery framework is based on recent work by the authors reported in Kumar et al. (2018c) and is sketched in Figure B.2.

B.3.1 *Long-Term MPC Formulation*

The elements of the long-term MPC formulation for the battery planning problem include the following model parameters, data, and variables:

Model Parameters and Data

- $L_{\xi,t} \in \mathbb{R}$: Buildings load [kW].
- $\pi^{e}_{\xi,t} \in \mathbb{R}$: Market price for electricity [\$/kWh].
- $\pi^{f}_{\xi,t} \in \mathbb{R}_{+}$: Market price for regulation capacity [\$/kW].
- $\pi^{D} \in \mathbb{R}_{+}$: Demand charge (monthly) [\$/kW].
- $\alpha_{\xi,t} \in [-1, 1]$: Fraction of FR capacity requested by ISO [-]. If $\alpha_t > 0$, the ISO sends a power to the battery while if $\alpha_t < 0$ the ISO withdraws power.
- $\overline{E} \in \mathbb{R}$: Battery storage capacity [kWh].
- $\overline{P} \in \mathbb{R}$: Maximum discharging rate (power) [kW].
- $\underline{P} \in \mathbb{R}$: Maximum charging rate (power) [kW].
- $\rho \in \mathbb{R}$: Minimum fraction of battery capacity reserved for frequency regulation [kWh/kW].
- $\overline{\Delta P} \in \mathbb{R}$: Maximum ramping limit [kW/h].

The forecasted disturbances in the battery problem are loads, prices, and regulation signals.

Model Variables

We define the following variables for the SP formulation, each replicated for all stages $\xi \in \Xi$:

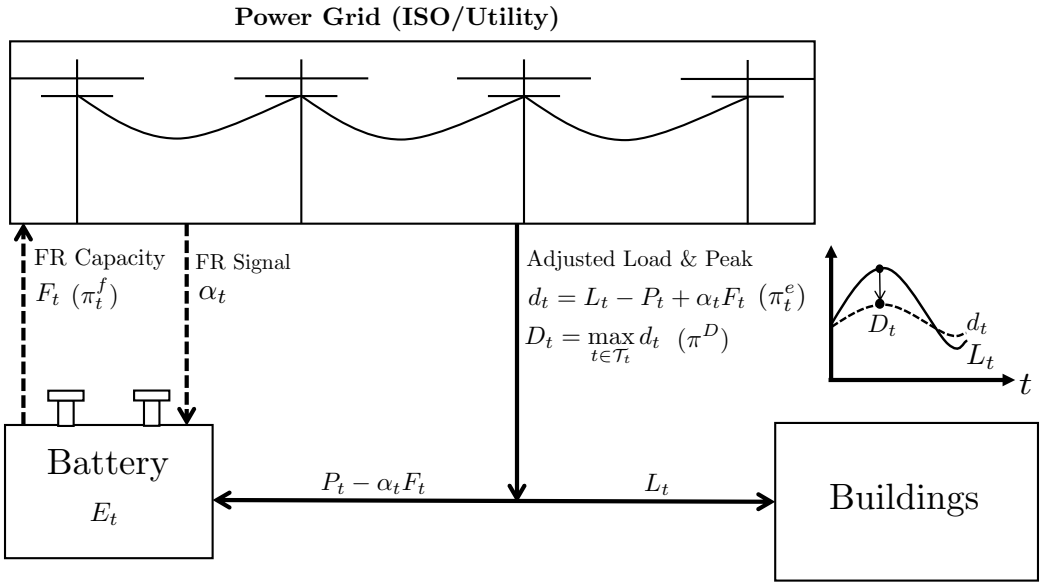

Figure B.2: Interactions battery, buildings, ISO, and utility.

- $P_{\xi,t} \in \mathbb{R}$: Net battery discharge rate (power) [kW].
- $F_{\xi,t} \in \mathbb{R}_+$: FR capacity provided to ISO [kW].
- $E_{\xi,t} \in \mathbb{R}_+$: State of charge (SOC) of the battery [kWh].
- $d_{\xi,t} \in \mathbb{R}_+$: Load requested from utility [kW].
- $D = \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} d_{\xi,t}$: Peak load over horizon $\mathcal{T}$ [kW]

Objective Function

We minimize total cost, which is given by the demand charge and the revenues collected from power and regulation:

$$\sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} \left(-\pi^e_{\xi,t}(P_{\xi,t} - \alpha_{\xi,t} F_{\xi,t}) - \pi^f_{\xi,t} F_{\xi,t} \right) + \pi^D D. \tag{B.3.6}$$

Here, the revenues represent the time-additive cost and the demand charge is the time-max cost.

Constraints

The constraints of the SP are replicated for every realization $\xi \in \Xi$. The net charged/discharged battery power plus the FR capacity provided must be within the maximum discharging and charging rates $\overline{P}$ and $\underline{P}$:

$$P_{\xi,t} + F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.7a}$$

$$P_{\xi,t} - F_{\xi,t} \geq -\underline{P},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.7b}$$

The storage dynamics are given by the difference equation:

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \bar{\mathcal{T}}_\xi, \xi \in \Xi \tag{B.3.8}$$

The following constraint is used to ensure that a certain amount of energy is reserved for the committed FR capacity over the interval $(t, t+1)$:

$$\rho F_{\xi,t} \leq E_{\xi,t} \leq \overline{E} - \rho F_{\xi,t},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.9a}$$

$$\rho F_{\xi,t} \leq E_{\xi,t+1} \leq \overline{E} - \rho F_{\xi,t},\ t \in \bar{\mathcal{T}}_\xi, \xi \in \Xi \tag{B.3.9b}$$

The battery ramp discharge rate is constrained as:

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}_\xi, \xi \in \Xi \tag{B.3.10}$$

The residual demand d_k requested from the utility is:

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.11}$$

The peak demand must satisfy $d_{\xi,t} \leq D,\ t \in \mathcal{T}_\xi, \xi \in \Xi$. We assume that the ISO does not allow the battery to sell back electricity. This is modeled by using the constraint $P_{\xi,t} + F_{\xi,t} \leq L_{\xi,t},\ t \in \mathcal{T}_\xi, \xi \in \Xi$. We enforce non-anticipativity constraint on the initial SOC $E_{\xi,0} = E_0$.

Finally, we enforce periodicity constraints on the SOC $E_{\xi,N_\xi} = E_0,\ \xi \in \Xi$. The bounds on the variables are:

$$0 \leq E_{\xi,t} \leq \overline{E},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.12a}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.12b}$$

$$0 \leq F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{B.3.12c}$$

The SP is solved to obtain the targets for the periodic SOC E_0^* and for the peak demand D^*. These targets are then used to guide a short-term MPC controller that we describe next.

B.3.2 *Short-Term MPC Formulation*

For simplicity, we assume that the short-term MPC controller only updates its control policies at the beginning of stage $t = t_\xi$ (where $t_\xi = \xi N_\xi,\ \xi \in \bar{\Xi}$) over horizon $\mathcal{T}_\xi := \{t, t+1, \ldots, t+N_\xi\}$. The short-term problem at time t_ξ uses forecasts for prices and loads over the horizon $\mathcal{T}_\xi$ (in the perfect information case this matches the scenarios of the long-term MPC formulation). The solution of the problem at time t_ξ is implemented for a block of N_ξ hours (i.e., N_ξ represents the update frequency in the MPC scheme). This approach is different than the traditional MPC approach in which the control policies are updated at every time step *within* the stage ξ. The short-term MPC formulation is:

$$\min_{P_{\xi,t},F_{\xi,t}} \sum_{t\in\mathcal{T}_\xi} \left(-\pi^e_{\xi,t}(P_{\xi,t} - \alpha_{\xi,t}F_{\xi,t}) - \pi^f_{\xi,t}F_{\xi,t}\right) + \pi^D D^* \tag{B.3.13a}$$

$$s.t.\quad P_{\xi,t} + F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi \tag{B.3.13b}$$

$$P_{\xi,t} - F_{\xi,t} \geq -\underline{P},\ t \in \mathcal{T}_\xi \tag{B.3.13c}$$

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t}F_{\xi,t},\ t \in \bar{\mathcal{T}}_\xi \tag{B.3.13d}$$

$$\rho F_{\xi,t} \leq E_{\xi,t} \leq \overline{E} - \rho F_{\xi,t},\ t \in \mathcal{T}_\xi \tag{B.3.13e}$$

$$\rho F_{\xi,t} \leq E_{\xi,t+1} \leq \overline{E} - \rho F_{\xi,t},\ t \in \bar{\mathcal{T}}_\xi \tag{B.3.13f}$$

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}_\xi \tag{B.3.13g}$$

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t} + \alpha_{\xi,t} F_{\xi,t},\ t \in \mathcal{T}_\xi \tag{B.3.13h}$$

$$P_{\xi,t} + F_{\xi,t} \leq L_{\xi,t},\ t \in \mathcal{T}_\xi \tag{B.3.13i}$$

$$E_{\xi,N} = E_0^* \tag{B.3.13j}$$

$$E_{\xi,0} = E_0^* \tag{B.3.13k}$$

$$d_{\xi,t} \leq D^* \tag{B.3.13l}$$

$$0 \leq E_{\xi,t} \leq \overline{E},\ t \in \mathcal{T}_\xi \tag{B.3.13m}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi \tag{B.3.13n}$$

$$0 \leq F_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi \tag{B.3.13o}$$

Here, E_0^* and D^* are fixed targets. In the perfect information case, the performance of the short-term MPC controller is optimal (because the true disturbance profiles match the forecasts used in the long-term MPC formulation). We now provide a strategy to update the short-term stage policies for the case of imperfect forecasts.

B.3.3 *Analysis for Imperfect Forecasts*

Under imperfect forecast case, the disturbance profile observed in real time can be different from the one that was used for long-term planning and determine the targets. Therefore, the target will not be optimal for the short-term MPC problem and can also be infeasible. To account for infeasibility, we add a penalty term to the cost function of the short-term MPC controller (B.3.13a) of the form $\pi^S S$. Here, the slack variable S is penalized by the cost π^S, where π^S must satisfy $\pi^S > \pi^D$ to guarantee that, if the peak target is feasible, then $S = 0$. In the case of an infeasible peak target, the MPC problem returns the smallest value of S, and therefore, the closest feasible operating point.

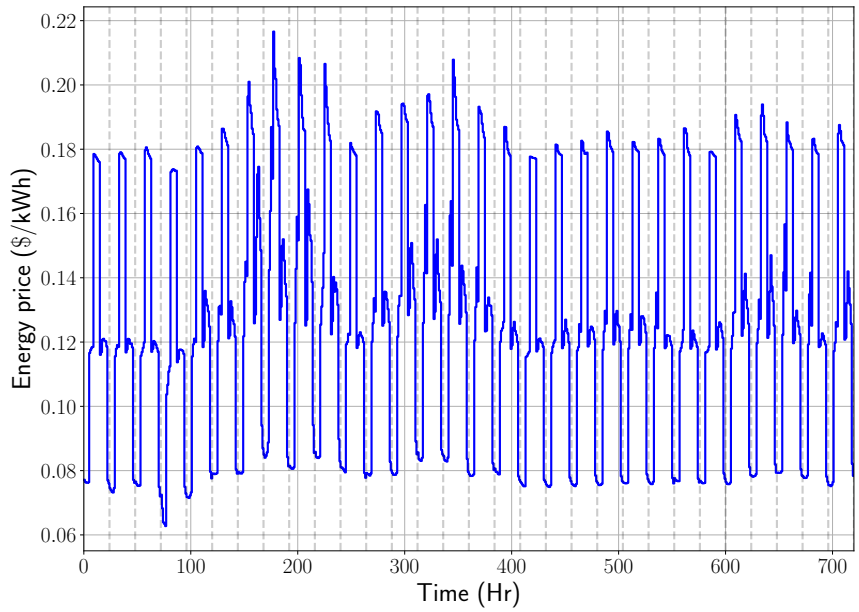

(a) Energy price for a month.

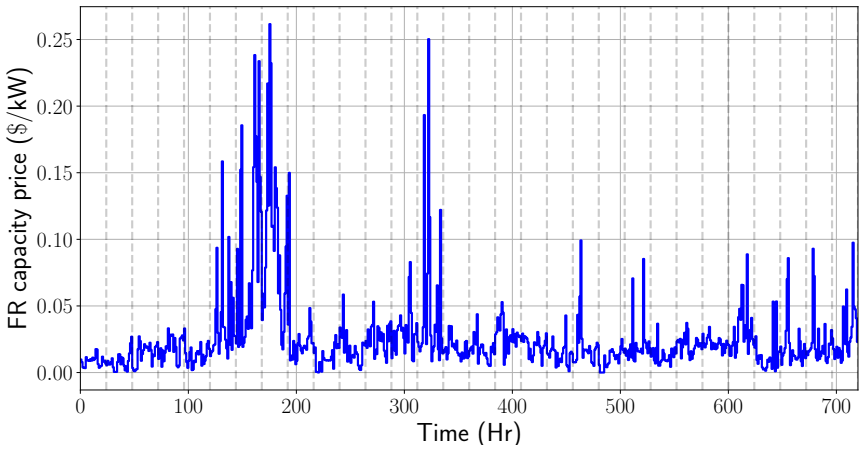

(b) Regulation capacity price for a month.

Figure B.3: Market price data used for the case studies.

B.4 Results

We consider a utility-scale stationary battery that has a capacity 0.5 MWh, rated power of 1 MW for both charge and discharge, and we assume a ramping limit of 0.5 MW/hr and at least 10% of battery capacity reserved for FR at every hour. We use historical data for one month for energy prices and FR prices from PJM Interconnection, shown in Figures B.3a, B.3b. Historical load data from a typical university campus for a month is used as the disturbance profile and is shown in the bottom of Figure B.4. We consider a planning horizon of one month (i.e., $N = 720$) and we use stages of 24 hours to create the SP formulation (i.e., $N_\xi = 24$ and $M = 30$).

We assume perfect forecasts in the first case study considered. For these experiments, we compare the cost of the hierarchical MPC scheme with the cost of a long-term MPC formulation *with and without periodicity constraints*. This comparison seeks to evaluate the impact of assuming an optimal periodic policy. In the second study, we evaluate the performance of the hierarchical MPC scheme under imperfect forecasts. To do so, we compare performance against a standard MPC scheme that performs hourly updates of the control policy and that uses a prediction horizon of 24 hours. For the standard MPC approach, we do not impose periodicity constraints. Instead, we use a discounting (weighting) factor of 1/30 for the demand charges (Kumar et al., 2018c).

Table B.1: Comparison of cost items under perfect forecasts.

Cost Item ($/month)	**Long-Term MPC** (with periodicity)	**Hierarhical MPC**	**Long-Term MPC** (without periodicity)
Total cost	114,079.81	114,079.81	114,077.99
Demand charge	129,424.86	129,424.86	129,424.86
FR	-14,861.72	-14,861.72	-14,863.76
Energy	-483.33	-483.33	-483.11

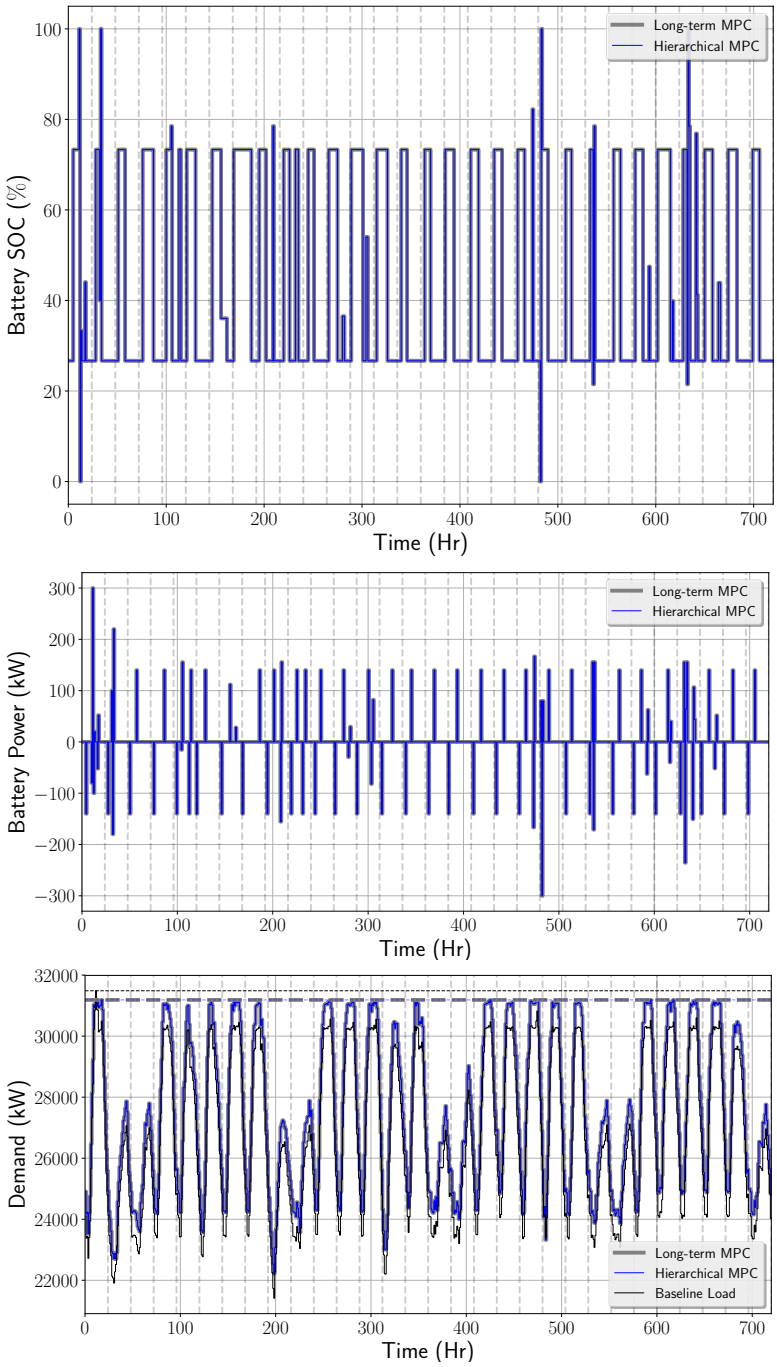

Figure B.4: Comparison of battery SOC (top), battery power (middle) and demand (bottom) policies.

B.4.1 *Perfect Forecasts*

Figure B.4 compares the policies obtained with the long-term MPC planning problem (labeled as Long-Term MPC) and the hierarchical MPC scheme (labeled as Hierarchical). The grey vertical lines denote 24-hour stages (scenarios). We see that the policies are identical; the equivalence indicates that the solutions of the stage subproblems are unique (for fixed targets D^* and E_0^*). Figure B.5 shows the SOC policy obtained from the long-term MPC problem with no periodicity constraints (B.2.1). It can be observed that the SOC is close to periodic. The cost items obtained with the different formulations are summarized in Table B.1. We see that the total cost obtained with no periodicity constraints is only 0.002% lower than that obtained with periodicity constraints (B.2.2). We thus conclude that assuming periodicity does not limit performance.

B.4.2 *Imperfect Forecasts*

Figure B.6 shows the forecasted load and the true (realized) load used in the case with imperfect forecasts. It can be observed that the peak in the realized load (33,315 kW) is 1800 kW higher than the peak in the forecasted load (31,486 kW). The design peak demand (D^*) obtained for the hierarchical MPC scheme using the forecasted load profile is 31,186 kW. This target value D^* leads to infeasibility of the short-term MPC controller when using the realized load profile. In Table B.2, we compare the cost items obtained with the different MPC schemes. We observe that the total cost obtained from the hierarchical MPC is 0.9% higher than that obtained with the long-term MPC formulation that uses the realized load to compute policies (perfect information). By comparing the demand charges of these two schemes, we see that hierarchical MPC is able to identify the optimal demand charge of long-term MPC. This is important because the demand charge is a significant component of the total cost. The higher total cost of hierarchical MPC is thus attributed to the suboptimal periodic SOC targets obtained from long-term MPC using the forecasted load profiles. We also evaluate the performance of standard MPC schemes, one that uses

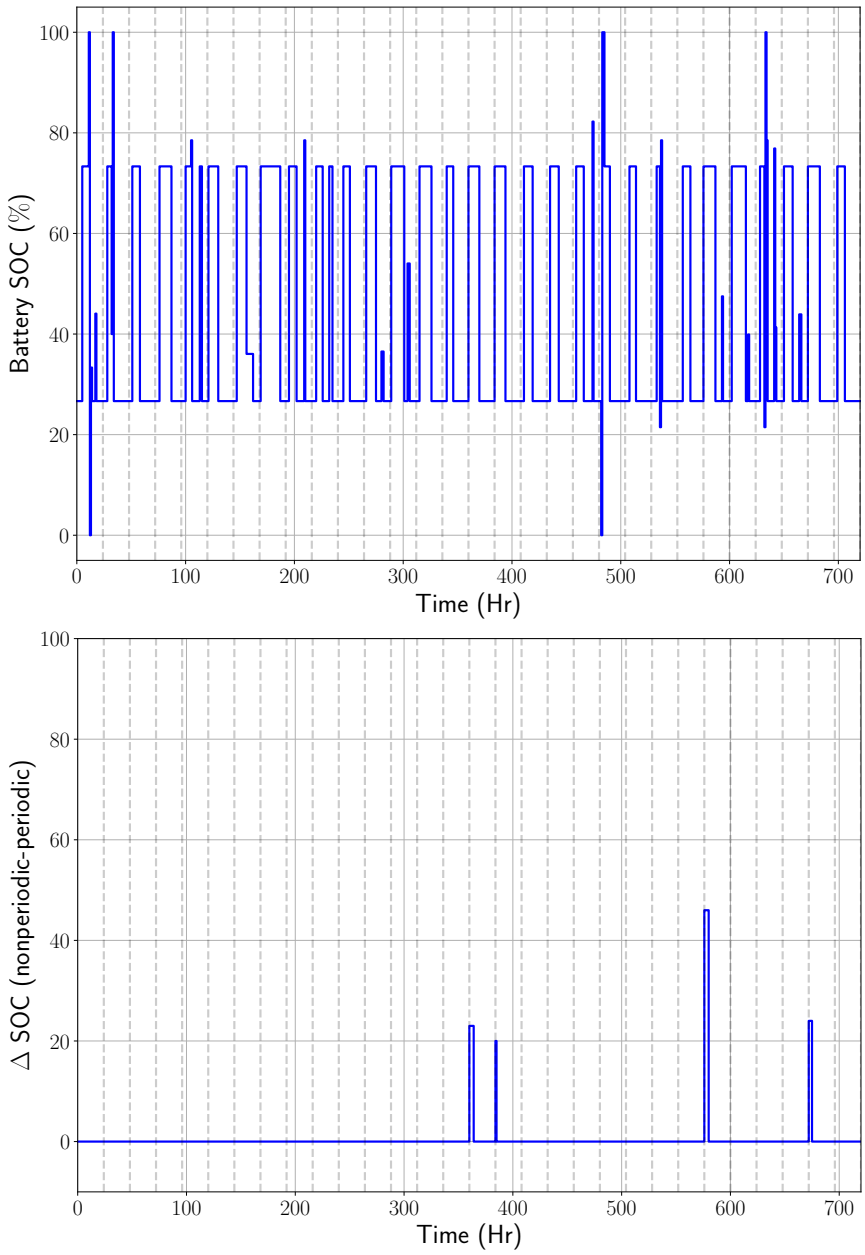

Figure B.5: SOC policy without periodicity (top). Difference of SOC policies with and without periodicity constraints (right).

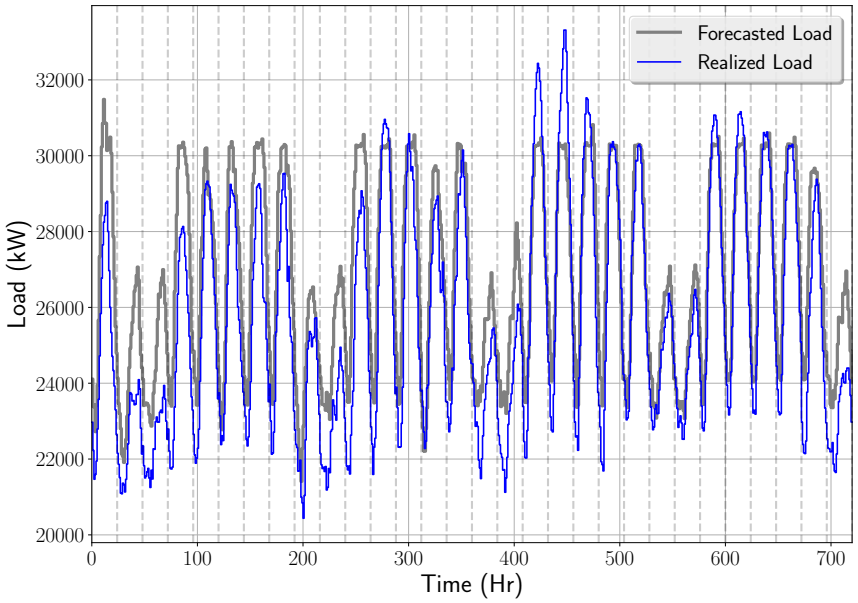

Figure B.6: Forecasted and realized load profiles.

the realized load (perfect information) and another that uses the forecasted load (imperfect information) to compute policies. By comparing long-term MPC and standard MPC with realized loads (perfect forecasts), we note that standard MPC yields a suboptimal policy. This highlights that the use of the discount factor only provides and ad-hoc approximation. Under imperfect forecasts, standard MPC results in a higher total cost compared to hierarchical MPC. This result highlights that hierarchical MPC provides a more effective approach to handle long-term demand charges because it systematically captures the load variability observed throughout the month. We also present a case study based on this work in Appendix C where we implement this hierarchical MPC scheme using stochastic programming for a battery system specifically for only mitigating demand charges, and show that this scheme is highly effective for applications where the major objective is the long-term peak cost.

Table B.2: Comparison of cost items under imperfect forecasts.

Cost Item ($/month)	**Long-Term MPC** (Forecasted Load)	**Long-Term MPC** (Realized Load)	**Hierarchical MPC** (Imperfect Forecast)	**Standard MPC** (Perfect Forecast)	**Standard MPC** (Imperfect Forecast)
Total cost	114,079.81	122,317.12	123,449.17	123,816.49	126,230.41
Demand charge	129,424.86	137,635.12	137,635.12	138,472.22	141,579.55
FR	-14,861.72	-14,850.11	-13,690.61	-14,178.68	-14,866.15
Energy	-483.33	-467.88	-495.33	-477.04	-482.98

B.5 Conclusions

We proposed an approach to handle long horizons in MPC based on the observation that if the periodicity constraints are enforced over short-term stages, the long horizon MPC problem can be cast as a stochastic programming (SP) problem. The SP setting reveals a mechanism to construct a hierarchical MPC scheme under which a high-level (long-horizon) MPC controller provides state targets to guide a low-level (short-horizon) MPC controller. We showed that this hierarchical MPC architecture is optimal under nominal (perfect forecast) conditions and can be extended to handle imperfect forecasts by correcting short-term policies. Another case study is presented in Appendix C based on this work where we implement this hierarchical MPC scheme using stochastic programming for a battery system only mitigating demand charges. Extensions to this work are presented in Chapter 5 that include updates (re-optimization) of the SP formulation to correct the periodic initial/terminal states. The use of an SP setting also opens the door to the use of scalable decomposition methods to tackle the long-term horizon problem. In particular, cutting-plane algorithms can be used to progressively update the high-level MPC layer as proposed in Chapter 5.

C

A HIERARCHICAL MODEL PREDICTIVE CONTROL APPROACH FOR HANDLING DEMAND CHARGES USING BATTERY SYSTEMS

In this appendix, we present another case study of the hierarchical MPC scheme designed in Chapter 5. In this case study, we focus specifically on mitigating long-term demand charge of a building system using stationary battery storage. We show that the hierarchical MPC scheme provides superior performance over standard MPC schemes that use ad-hoc approaches to handle the multiple timescales arising from long-term demand charges and short-term energy prices and loads.

C.1 Introduction

Peak electricity demands or the demand charges remain a serious concern for utilities since they pose capacity challenges to the power grid. This situation has provided opportunities to energy storage systems to create savings by reducing the peak electricity demands of the buildings and campuses (de Salis et al., 2014; Johnson et al., 2011; Oudalov et al., 2007a; Rahimi et al., 2013; Lu et al., 2014; Dong et al., 2011; Joshi and Pindoriya, 2015). In this context, MPC is used in Braun (1990); Ma et al. (2012a) to mitigate demand charges of a building HVAC system. Demand charges pose a challenging multiscale planning prob-

lem that should make decisions at fine timescales while mitigating long-term costs. In particular, handling MPC planning formulations over long horizons can be intractable. To overcome this issue, the authors in Braun (1990); Ma et al. (2012a) penalize the peak electricity demand over the short-term receding horizon. This approach is practical but can yield overly conservative policies and deteriorate economic performance. Long-term discounting cost factors are used in Risbeck et al. (2017); Patel et al. (2016) to reduce the conservatism of short-term MPC formulations. This approach is intuitive but does not provide optimality guarantees. The work in Zavala (2016) proposes a hierarchical MPC scheme, with the adjoint (dual) information obtained from a long-term and coarse MPC controller used to guide the control policy of a short-term MPC controller. It is demonstrated that this approach can achieve optimality. Unfortunately, continuity of the adjoint profiles is necessary for this approach, which is not guaranteed in general. Alternative hierarchical MPC schemes (Scattolini and Colaneri, 2007; Picasso et al., 2010) available in the literature provide feasibility but not optimality guarantees. Moreover, such approaches cannot handle peak costs.

In this work, the demand charges for buildings are handled by battery systems by using a recently developed approach for hierarchical model predictive control (MPC). Here, a long-term planner MPC provides guiding targets to the short-term MPC controller. In the proposed hierarchical MPC approach, the state of charge (SOC) policy is assumed to be periodic, which is a reasonable assumption in the context of energy systems because loads and price profiles have strong periodic components (Huang et al., 2011; Subramanian et al., 2014; Risbeck et al., 2015). This assumption makes it possible to cast the long-term planning problem as a tractable stochastic programming problem. Here, the period (e.g., a day or week) represents an operational scenario and we seek to determine targets for the periodic SOC levels and the peak cost. The periodic SOC targets and peak cost are communicated to a short-term MPC controller. The intra-period charge/discharge policies (at high resolution) are determined by the short-term MPC controller while meeting the targets obtained from the long-term planning.

The performance of the hierarchical MPC scheme is demonstrated using an application in buildings with electricity storage through simulation case studies. Here, the goal is to determine the optimal short-term (hourly) charge-discharge policy while mitigating long-term (monthly) demand charges from utilities. Using the simulation case studies, it is shown that the proposed hierarchical MPC scheme yields optimal demand charge and charge-discharge policy under nominal (perfect forecast) conditions. Comparative studies of the proposed hierarchical MPC scheme and standard MPC schemes that use ad-hoc approaches to handle the multiple timescales are also presented. The solution of the hierarchical MPC scheme is compared with that from a standard MPC scheme that uses a discounting (weighting) factor to account for the inability of solving the long horizon problem. Under imperfect forecasts, it is observed that the cost from the hierarchical MPC scheme provides superior performance over standard MPC because it can capture long-term variability in a more systematic manner.

C.2 Problem Formulation

We derive the problem forumulation based on the work in Appendix B. We begin by considering the following MPC problem:

$$\min_{\mathbf{x}_{\mathcal{N}},\mathbf{u}_{\mathcal{N}}} \sum_{t=0}^{N-1} a_t^T x_t + b_t^T u_t + c_t^T d_t + \max_{t\in\mathcal{T}}(p_t^T x_t + q_t^T u_t + r_t^T d_t) \tag{C.2.1a}$$

$$\text{s.t. } x_{t+1} = Ax_t + Bu_t + Cd_t,\ t = 0,\dots,N-1 \tag{C.2.1b}$$

$$u_t \in \mathcal{U},\ t = 0,\dots,N-1 \tag{C.2.1c}$$

$$x_t \in \mathcal{X},\ t = 0,\dots,N \tag{C.2.1d}$$

Here, $a_t \in \mathbb{R}^{n_x}$, $b_t \in \mathbb{R}^{n_u}$ and $c_t \in \mathbb{R}^{n_u}$ are time-additive stage costs associated with the states $x_t \in \mathbb{R}^{n_x}$, controls $u_t \in \mathbb{R}^{n_u}$, and $d_t \in \mathbb{R}^{n_d}$, respectively. $p_t \in \mathbb{R}^{n_x}$, $q_t \in \mathbb{R}^{n_u}$ and $r_t \in \mathbb{R}^{n_u}$ are the costs associated with the states $x_t \in \mathbb{R}^{n_x}$, controls $u_t \in \mathbb{R}^{n_u}$, and $d_t \in \mathbb{R}^{n_d}$, respectively, of which only the peak cost (time-max cost) is considered. The horizon

length is denoted as N. Trajectories for control actions and states over the horizon are denoted as $\mathbf{u}_{\mathcal{N}} := (u_0, u_1, \ldots, u_{N-1})$, and $\mathbf{x}_{\mathcal{N}} := (x_1, x_2, \ldots, x_N)$. The dynamics of the linear system are described by the matrices $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$, and $C \in \mathbb{R}^{n_x \times n_d}$. The state x_0 is provided as the initial condition. The states x_t and controls u_t are bounded by the polyhedral sets $\mathcal{X}$ and $\mathcal{U}$, respectively.

The time set is represented by $\mathcal{T} := \{0, 1, \ldots, N\}$ and it is also considered that the set $\mathcal{T}$ is partitioned (in lexicographic order) into a set of short time *periods* $\Xi := \{0, \ldots, M\}$ with each period $\xi \in \Xi$ comprising of equal number of time steps N_T satisfying $M \times N_T = N$. Further, each period is defined as $\mathcal{T}_\xi := \{0, \ldots, N_T\}$. For convenience, the sets $\bar{\Xi} := \Xi \setminus \{M\}$ and $\bar{\mathcal{T}}_\xi := \mathcal{T}_\xi \setminus \{N_T\}$ are also defined. We partition states, controls, and disturbance policies into stages and denote the stage policies as $u_{\xi,t}$, $x_{\xi,t}$, and $d_{\xi,t}$ for $\xi \in \Xi$ and $t \in \mathcal{T}_\xi$. For compactness in notation, $u_{\xi,N_T} = 0$ and $d_{\xi,N_T} = 0$ are defined for all periods ξ. These partitions are used to reformulate the MPC problem in the following *equivalent* form:

$$\min_{u_{\xi,t}} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} a_{\xi,t}^T x_{\xi,t} + b_{\xi,t}^T u_{\xi,t} + c_{\xi,t}^T d_{\xi,t} + \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} (p_{\xi,t}^T x_{\xi,t} + q_{\xi,t}^T u_{\xi,t} + r_{\xi,t}^T d_{\xi,t}) \quad \text{(C.2.2a)}$$

$$\text{s.t. } x_{\xi,t+1} = Ax_{\xi,t} + Bu_{\xi,t} + Cd_{\xi,t},\ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \quad \text{(C.2.2b)}$$

$$x_{\xi+1,0} = x_{\xi,N_T}, \xi \in \bar{\Xi} \quad \text{(C.2.2c)}$$

$$x_{0,0} = \bar{x}_0 \quad \text{(C.2.2d)}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \quad \text{(C.2.2e)}$$

Here, the constraint (C.2.2c) enforces continuity between stages.

By enforcing periodicity of the state at the end of every stage, the following modified problem is obtained:

$$\min_{u_{\xi,t}} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} a_{\xi,t}^T x_{\xi,t} + b_{\xi,t}^T u_{\xi,t} + c_{\xi,t}^T d_{\xi,t} + \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} (p_{\xi,t}^T x_{\xi,t} + q_{\xi,t}^T u_{\xi,t} + r_{\xi,t}^T d_{\xi,t}) \quad \text{(C.2.3a)}$$

$$\text{s.t. } x_{\xi,t+1} = Ax_{\xi,t} + Bu_{\xi,t} + Cd_{\xi,t},\ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \quad \text{(C.2.3b)}$$

$$x_{\xi+1,0} = x_{\xi,N_T}, \xi \in \bar{\Xi} \quad \text{(C.2.3c)}$$

$$x_{\xi,N_T} = x_{\xi,0}, \xi \in \Xi \tag{C.2.3d}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{C.2.3e}$$

The optimization formulation (C.2.3) represents the *l*ong-term MPC planning problem with the enforced periodicity constraints. Here, the variable $x_{0,0}$ is a free variable that is sought to be optimized. By combining the periodicity constraint (C.2.3d) and the stage continuity constraints (C.2.3c), a reformulated stage continuity constraints can be obtained in the form of $x_{\xi+1,0} = x_{\xi,0}$, $\xi \in \bar{\Xi}$. A *lifting* variable x_0 is further introduced to reformulate $x_{\xi+1,0} = x_{\xi,0}$, $\xi \in \bar{\Xi}$ as $x_{\xi,0} = x_0$, $\xi \in \Xi$. Consequently, the goal of the long-term MPC formulation (C.2.3) is to find the optimal periodic state x_0 and control policies $u_{\xi,t}$, $\xi \in \Xi, t \in \mathcal{T}_\xi$ that minimize the time-additive and peak costs. We also reformulate the peak cost function to obtain the following final equivalent form of (C.2.3):

$$\min_{u_{\xi,t}} \sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} a_{\xi,t}^T x_{\xi,t} + b_{\xi,t}^T u_{\xi,t} + c_{\xi,t}^T d_{\xi,t} + \eta \tag{C.2.4a}$$

$$\text{s.t.} p_{\xi,t}^T x_{\xi,t} + q_{\xi,t}^T u_{\xi,t} + r_{\xi,t}^T d_{\xi,t} \leq \eta, \ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \tag{C.2.4b}$$

$$x_{\xi,t+1} = A x_{\xi,t} + B u_{\xi,t} + C d_{\xi,t}, \ \xi \in \Xi, t \in \bar{\mathcal{T}}_\xi \tag{C.2.4c}$$

$$x_{\xi+1,0} = x_0, \xi \in \bar{\Xi} \tag{C.2.4d}$$

$$x_{\xi,N_T} = x_0, \xi \in \Xi \tag{C.2.4e}$$

$$x_{\xi,t} \in \mathcal{X}, u_{\xi,t} \in \mathcal{U}. \tag{C.2.4f}$$

The solution of the optimization problem (C.2.4) is denoted as $x_{\xi,t}^*, u_{\xi,t}^*, \eta^*$, where η^* is the peak cost over the entire planning horizon. It is also noted that $x_{\xi+1,0}^* = x_{\xi,0}^* = x_0^*$.

The structure of (C.2.4) reveals that the only coupling variables between stages are x_0 and η. Therefore, (C.2.4) can be seen as a *stochastic programming problem* in which periods are operational scenarios, x_0 and η are *design variables* or *first-stage variables*, and $x_{\xi,t}, u_{\xi,t}$ are *recourse policies* for scenarios. The problem (C.2.4) is decomposed into M subproblems

by fixing the design variables to their optimal values x_0^* and η^*:

$$\min_{u_{\xi,t}} \sum_{\xi\in\Xi}\sum_{t\in\mathcal{T}_\xi} a_{\xi,t}^T x_{\xi,t} + b_{\xi,t}^T u_{\xi,t} + c_{\xi,t}^T d_{\xi,t} + \eta^* \tag{C.2.5a}$$

$$\text{s.t.} p_{\xi,t}^T x_{\xi,t} + q_{\xi,t}^T u_{\xi,t} + r_{\xi,t}^T d_{\xi,t} \le \eta^*,\ \xi\in\Xi, t\in\bar{\mathcal{T}}_\xi \tag{C.2.5b}$$

$$x_{\xi,t+1} = Ax_{\xi,t} + Bu_{\xi,t} + Cd_{\xi,t},\ \xi\in\Xi, t\in\bar{\mathcal{T}}_\xi \tag{C.2.5c}$$

$$x_{\xi+1,0} = x_0^*, \xi\in\bar{\Xi} \tag{C.2.5d}$$

$$x_{\xi,N_T} = x_0^*, \xi\in\Xi \tag{C.2.5e}$$

$$x_{\xi,t}\in\mathcal{X}, u_{\xi,t}\in\mathcal{U}. \tag{C.2.5f}$$

It is noted that the subproblem (C.2.5) has the structure of a standard MPC problem with periodicity constraints. The SP formulation (C.2.4) provides an opportunity to implement a *hierarchical MPC scheme*, in which the long-term MPC problem of the form (C.2.4) (equivalently (C.2.3)) can guide a short-term MPC controller of the form (C.2.4), with the targets for the design variables obtained from the long-term MPC. The targets for the periodic state x_0^* and the peak cost η^* provide a form of communication between the hierarchical levels. Figure C.1 provides a schematic representation of the hierarchical MPC scheme.

From the SP setting, it is also revealed that when the *disturbance forecasts are perfect*, the solution of the short-term MPC controller will yield an optimal stage trajectory $u_{\xi,t}^*$ and $x_{\xi,t}^*$ since the targets $x_{\xi,0}^*$ and η^* are optimal, and these targets also correspond to the minimum cost for the entire horizon (as obtained from the long-term MPC problem). The SP setting also indicates that, when the forecasts are imperfect, the targets x_0^*, η^* may not be optimal and it may not be feasible for the subproblem (C.2.5) to achieve the provided design targets by the long-term MPC planner. This situation can be mitigated by penalizing deviations from the target (which will find the closest feasible point) or by re-optimizing the design targets using the actual realized disturbances when they become available in real-time. The SP setting also reveals that it is possible to find targets that remain optimal and feasible for all realizations (or many realizations) by considering a larger number of sce-

narios of the disturbance in the planning problem. Consequently, the proposed approach provides a framework to easily construct robust formulations. Another important feature of the hierarchical MPC approach is that there is no need to weight or discount the long-term costs in the short-horizon MPC subproblem because the long-term effects in the costs are already accounted for by the long-term MPC. Moreover, there already exist advanced techniques to solve the SP problem efficiently by using decomposition schemes based on parallel linear algebra and Benders/Lagrangian decomposition (Zavala et al., 2008; CarøE and Schultz, 1999; Geoffrion, 1972).

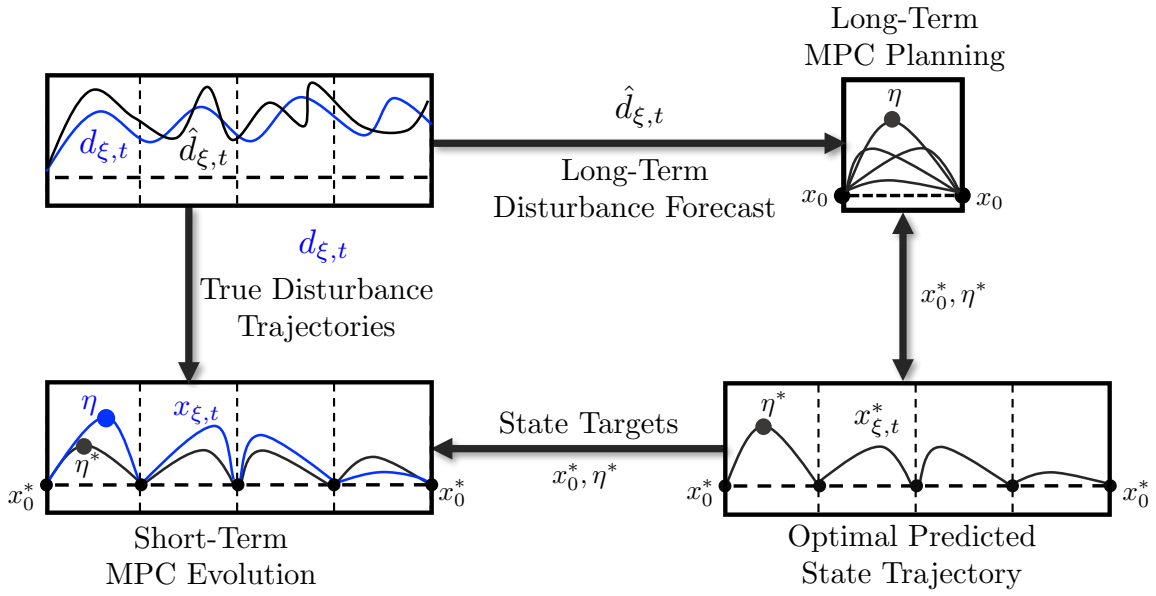

Figure C.1: Hierarchical MPC scheme.

C.3 Stationary Battery Case Study

Batteries are flexible assets that can be used to provide energy and can aid utility companies by providing demand-side management capabilities for buildings or manufacturing facilities (Rastler, 2010; Oudalov et al., 2006). The use of batteries for demand charge mitigation has been studied in Kumar et al. (2018a); Sigrist et al. (2013); White and Zhang (2011); Lucas and Chondrogiannis (2016); Sebastián (2016); Oudalov et al. (2006); Shi et al.

(2017). In such settings, the objective function of the MPC scheme takes the form of the formulation in (C.2.1), where the additive costs represent the total time-of-use energy cost and peak costs represent demand charges. The setting considered in this work is sketched in Figure C.2.

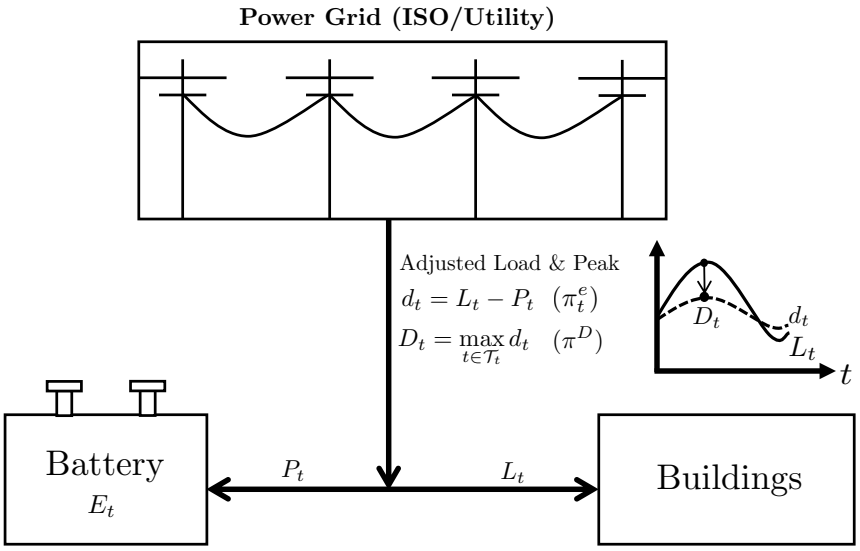

Figure C.2: Interactions between battery, buildings, ISO, and utility.

C.3.1 *Long-Term MPC Formulation*

The elements of the long-term MPC formulation for the battery planning problem include the model parameters, data, and variables, which are listed below.

Model Parameters and Data

- $L_{\xi,t} \in \mathbb{R}$: Buildings load [kW].
- $\pi^e_{\xi,t} \in \mathbb{R}$: Market price for electricity [$/kWh].
- $\pi^D \in \mathbb{R}_+$: Demand charge (monthly) [$/kW].
- $\overline{E} \in \mathbb{R}$: Battery storage capacity [kWh].
- $\overline{P} \in \mathbb{R}$: Maximum discharging rate (power) [kW].

- $\underline{P} \in \mathbb{R}$: Maximum charging rate (power) [kW].
- $\overline{\Delta P} \in \mathbb{R}$: Maximum ramping limit [kW/h].

Model Variables

- $P_{\xi,t} \in \mathbb{R}$: Net battery discharge rate (power) [kW]. If $P_{\xi,t} > 0$, the battery is being discharged and if $P_{\xi,t} < 0$ the battery is being charged.
- $E_{\xi,t} \in \mathbb{R}_+$: State of charge (SOC) of the battery [kWh].
- $d_{\xi,t} \in \mathbb{R}_+$: Load requested from utility [kW].
- $D = \max_{\xi \in \Xi} \max_{t \in \mathcal{T}_\xi} d_{\xi,t}$: Peak load over horizon $\mathcal{T}$ [kW]

Objective Function

The objective is to minimize total cost, which is given by the demand charge and the revenues collected from power and regulation:

$$\sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} \pi^e_{\xi,t}(L_{\xi,t} - P_{\xi,t}) + \pi^D D. \tag{C.3.6}$$

Here, the energy cost $\pi^e_{\xi,t}(L_{\xi,t} - P_{\xi,t})$ represents the time-additive cost and the demand charge is the time-max cost. The component $\pi^e_{\xi,t} L_{\xi,t}$ in the cost function is just a constant with respect to the optimization, and therefore we can drop this component from the cost function. The component $-\pi^e_{\xi,t} P_{\xi,t}$ provides the cost saving by using the battery energy instead of buying energy directly from the utility. Therefore, the following cost function is considered for the further analysis:

$$\sum_{\xi \in \Xi} \sum_{t \in \mathcal{T}_\xi} -\pi^e_{\xi,t} P_{\xi,t} + \pi^D D. \tag{C.3.7}$$

Constraints

The constraints on the system are the physical charging/discharging limits of the battery, the battery state of charge (SOC) dynamics, and the peak demand computation. The constraints of the SP are replicated for every scenario (or period) $\xi \in \Xi$. The storage dynamics are given by the difference equation:

$$E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t},\ t \in \bar{\mathcal{T}}_{\xi}, \xi \in \Xi \tag{C.3.8}$$

The battery ramp discharge rate is constrained as:

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}_{\xi}, \xi \in \Xi \tag{C.3.9}$$

The residual demand d_k requested from the utility is:

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \tag{C.3.10}$$

The peak demand must satisfy:

$$d_{\xi,t} \leq D,\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \tag{C.3.11}$$

It is assumed that the ISO does not allow the battery to sell back electricity. This is modeled by using the constraint:

$$P_{\xi,t} \leq L_{\xi,t},\ t \in \mathcal{T}_{\xi}, \xi \in \Xi \tag{C.3.12}$$

The initial SOC is is a design variable which is enforced by using the following non-anticipativity constraint:

$$E_{\xi,0} = E_0,\ \xi \in \Xi \tag{C.3.13}$$

Finally, the periodicity constraints are enforced, i.e. the final state of charge in each scenario is the same as the initial state:

$$E_{\xi,N_T} = E_0,\ \xi \in \Xi \tag{C.3.14}$$

The bounds on the variables are given by:

$$0 \leq E_{\xi,t} \leq \overline{E},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{C.3.15a}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi, \xi \in \Xi \tag{C.3.15b}$$

The SP is solved to obtain the targets for the periodic battery SOC E_0^* and for the peak demand D^*. These targets are then used to guide a short-term MPC controller that obtains the battery operating policy at every stage.

C.3.2 *Short-Term MPC Formulation*

For simplicity in the presentation, it is assumed that the short-term MPC controller only updates its control policies at the beginning of every stage $t = t_\xi$ (where $t_\xi = \xi N_T$, $\xi \in \bar{\Xi}$) over horizon $\mathcal{T}_\xi := \{t, t+1, \ldots, t+N_T\}$. The short-term problem at time t_ξ uses forecasts for prices and loads over the prediction horizon $\mathcal{T}_\xi$ (in the perfect information case this matches the scenarios of the long-term MPC formulation). The solution of the problem at time t_ξ is implemented for a block of N_T hours. This approach is different from the traditional approach in which the control policies are updated at every time step *within* the stage ξ. The short-term MPC formulation in stage ξ is given by:

$$\min_{P_{\xi,t}, F_{\xi,t}} \sum_{t \in \mathcal{T}_\xi} \pi^e_{\xi,t}(L_{\xi,t} - P_{\xi,t}) + \pi^D D^* \tag{C.3.16a}$$

$$s.t. \quad E_{\xi,t+1} = E_{\xi,t} - P_{\xi,t},\ t \in \bar{\mathcal{T}}_\xi \tag{C.3.16b}$$

$$-\overline{\Delta P} \leq P_{\xi,t+1} - P_{\xi,t} \leq \overline{\Delta P},\ t \in \bar{\mathcal{T}}_\xi \tag{C.3.16c}$$

$$d_{\xi,t} = L_{\xi,t} - P_{\xi,t},\ t \in \mathcal{T}_\xi \tag{C.3.16d}$$

$$P_{\xi,t} \leq L_{\xi,t},\ t \in \mathcal{T}_\xi \tag{C.3.16e}$$

$$E_{\xi,N} = E_0^* \tag{C.3.16f}$$

$$E_{\xi,0} = E_0^* \tag{C.3.16g}$$

$$d_{\xi,t} \leq D^* \tag{C.3.16h}$$

$$0 \leq E_{\xi,t} \leq \overline{E},\ t \in \mathcal{T}_\xi \tag{C.3.16i}$$

$$-\underline{P} \leq P_{\xi,t} \leq \overline{P},\ t \in \mathcal{T}_\xi \tag{C.3.16j}$$

C.4 Results

In this section, the results from the simulation case studies based on the MPC formulation described in Section C.3 are presented. For the simulation case studies, a utility-scale stationary battery that has a capacity 0.5 MWh and rated power of 1 MW for both charge and discharge is considered, and a ramping limit of 0.5 MW/hr is assumed. Historical data for one month for energy prices from PJM Interconnection, shown in Figures C.3a are used (PJM is a regional transmission organization (RTO) that coordinates the movement of wholesale electricity in all or parts of 13 states and the District of Columbia in the United States). Historical load data from a typical university campus for a month is used as the disturbance profile shown in Figure C.3b. From Figures C.3, it can be clearly observed the disturbance profiles have strong periodic components and therefore, the proposed hierarchical MPC scheme based on enforcing periodicity is an appropriate approach for planning of such systems. A planning horizon of one month (i.e., $N = 720$) is considered and stages of 24 hours are used to create the SP formulation (i.e., $N_T = 24$ and $M = 30$).

Perfect forecasts are assumed in the first case study considered. In these experiments, the cost of the hierarchical MPC scheme is compared with the cost of a long-term MPC formulation *with and without periodicity constraints*. This comparison seeks to evaluate the impact of assuming an optimal periodic policy. In the second study, the performance of the

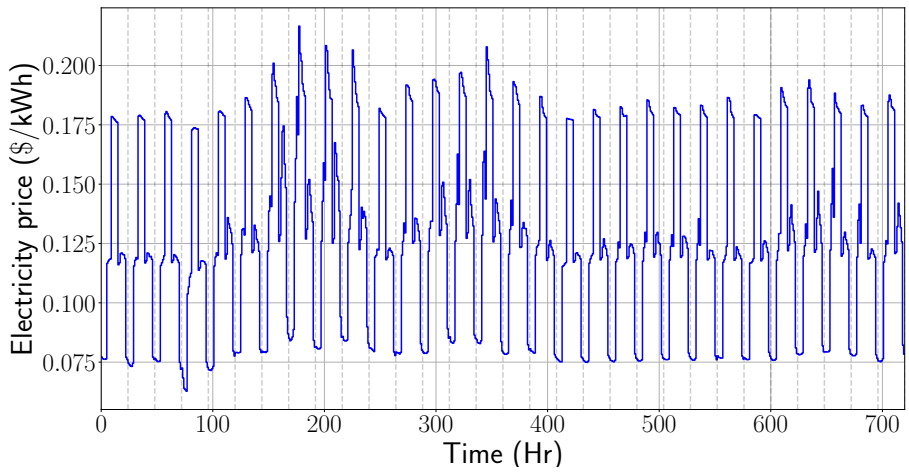

(a) Market electricity for a month.

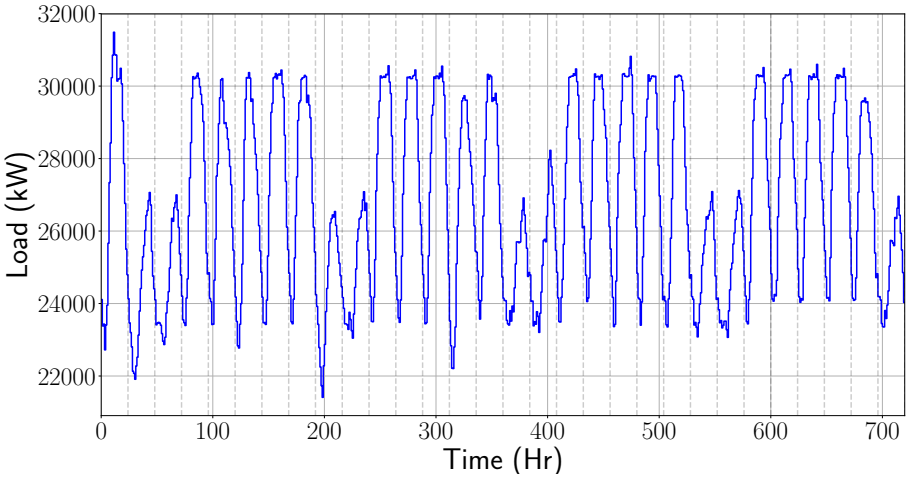

(b) University campus load for a month.

Figure C.3: Market and load data used for the case studies.

hierarchical MPC scheme is evaluated under imperfect forecasts. To do so, its performance is compared against a standard MPC scheme that performs hourly updates of the control policy and that uses a prediction horizon of 24 hours. For the standard MPC approach, periodicity constraints are not imposed. Instead, a discounting (weighting) factor is used for the demand charges based on the length of horizon. In this case study, a discounting factor of $\frac{1}{30}$ is used.

C.4.1 *Perfect Forecasts*

Figures C.4-C.5b compare the policies obtained with the long-term MPC planning problem and the hierarchical MPC scheme. The grey vertical lines denote 24-hour periods (scenarios). It can be seen that the policies are identical; the equivalence indicates that the solutions of the period subproblems are unique (for fixed targets D^* and E_0^*).

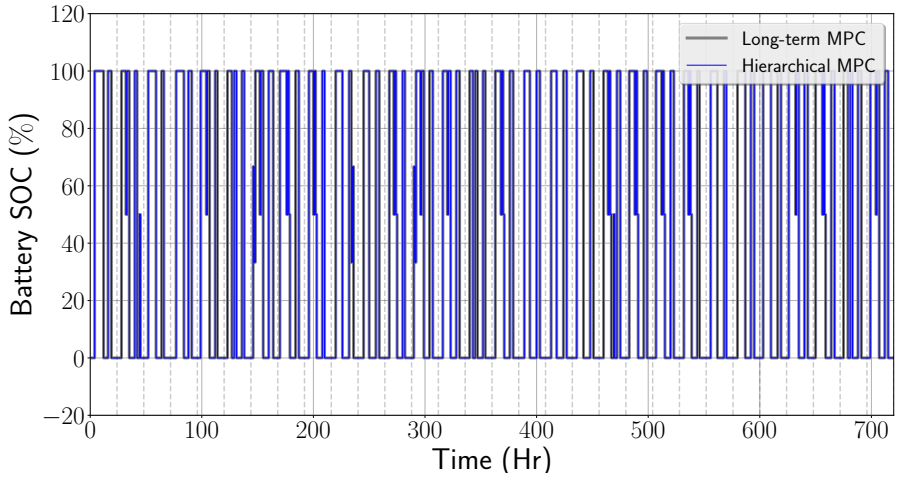

(a) Battery SOC trajectories.

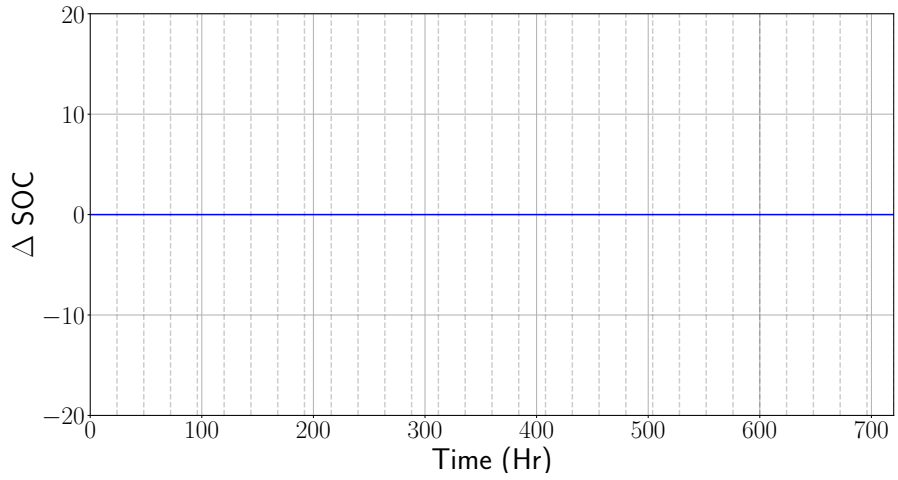

(b) Difference of battery SOC trajectories.

Figure C.4: Comparison of SOC policies.

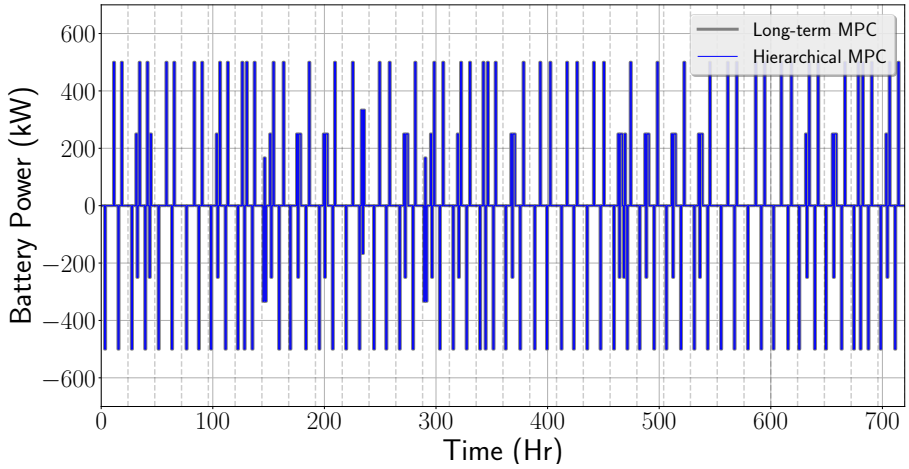

(a) Battery discharge (power) policies.

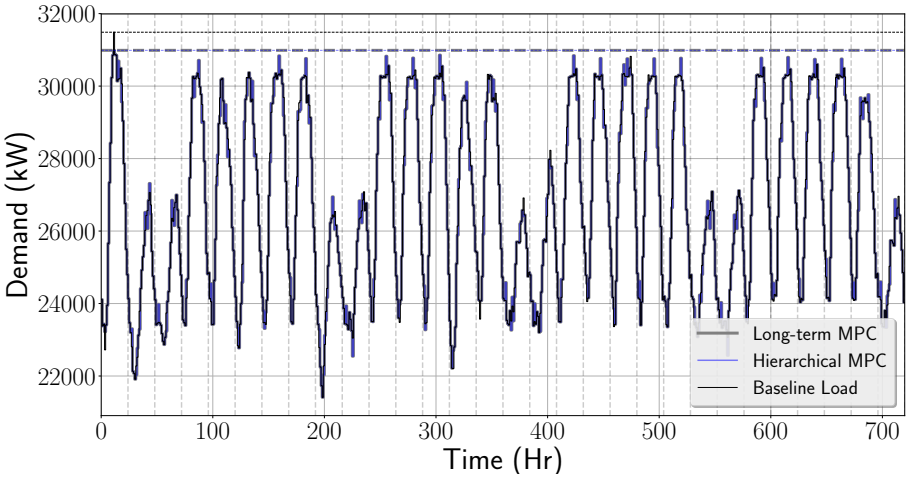

(b) Comparison of demand policies.

Figure C.5: Comparison of battery power and building demand policies.

Table C.1: Comparison of cost items under perfect forecasts.

Cost Item ($/month)	**Long-Term MPC** (with periodicity)	**Hierarhical MPC**	**Long-Term MPC** (without periodicity)
Total cost	126,690.05	126,690.05	126,637.27
Demand charge	128,594.86	128,594.86	128,594.86
Energy	-1,904.81	-1,904.81	-1,957.59

Figure C.6 shows the SOC policy obtained from the long-term MPC problem (C.2.2) without the periodicity constraints. It can be observed that the SOC is quasi-periodic. The

cost items obtained with the different formulations are summarized in Table C.1. It is also found that the total cost obtained with no periodicity constraints is only 0.04% lower than that obtained with periodicity constraints (C.2.3). It can thus be concluded that assuming a periodic policy does not limit performance.

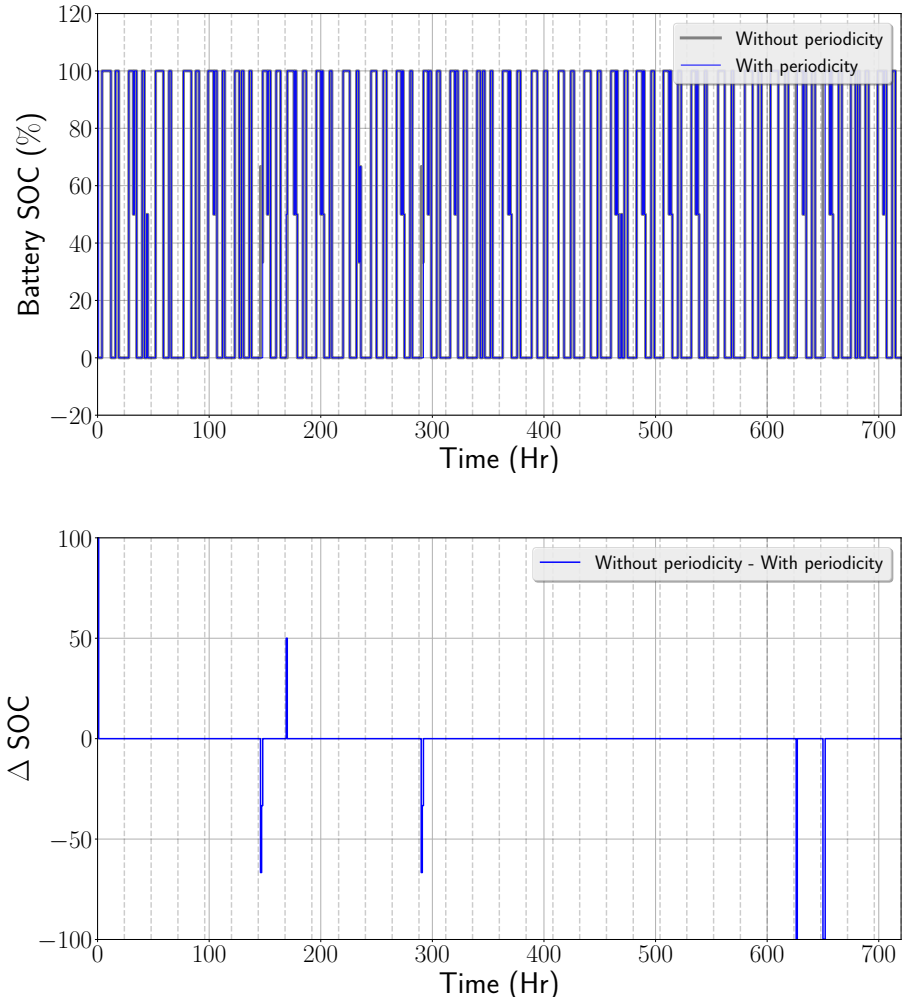

Figure C.6: SOC policy without periodicity (left). Comparison of SOC policies with and without periodicity (right).

C.4.2 *Imperfect Forecasts*

Figure C.7 shows the forecasted load and the realized load used in the case studies with imperfect forecasts. It can be observed that the peak in the realized load (33,315.55 kW) is more than 1800 kW higher than the peak in the forecasted load (31,486.71 kW). The design peak demand (D^*) obtained for the hierarchical MPC scheme using the forecasted load profile is 31,186.71 kW. This target value D^* leads to infeasibility of the short-term MPC controller when using the realized load profile because the size and power rating of the battery is not able to achieve this provided target for peak demand.

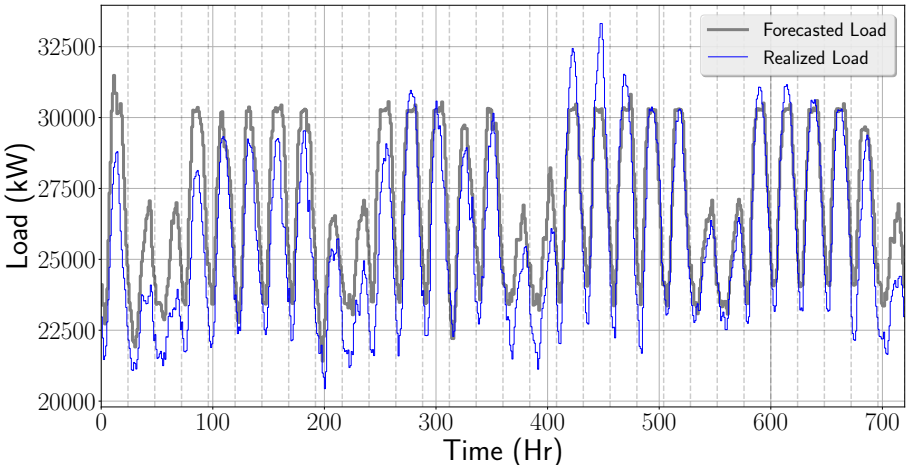

Figure C.7: Forecasted and realized load profiles.

In Table C.2, the cost items obtained with the different MPC schemes are compared. It is observed that the total cost obtained from the hierarchical MPC is only 0.02% higher than that obtained with the long-term MPC formulation that uses the advanced knowledge of the realized load to compute policies (perfect information). It is found that hierarchical MPC is able to identify the optimal demand charge of long-term MPC even under imperfect forecast. This is important because the demand charge is a significant component of the total cost. The higher total cost of hierarchical MPC is thus attributed to the suboptimal periodic SOC levels obtained from the long-term MPC with the forecasted load profiles.

The performance of standard MPC schemes, one that uses the realized load (perfect information) and another that uses the forecasted load (imperfect information) to compute policies are also evaluated. By comparing long-term MPC and standard MPC with realized loads (perfect forecasts), it can be immediately noted that standard MPC yields a suboptimal policy and a higher total cost. This highlights that the use of the discount factor only provides an ad-hoc approximation. Under imperfect forecasts, standard MPC also results in 2.24% higher total cost compared to hierarchical MPC. These results highlight the lack of robustness of the standard MPC approach. This also highlights that even under the imperfect forecast case, the hierarchical MPC provides a more effective approach to handle long-term demand charges because it systematically captures the load variability observed throughout the month and handles the long-term cost effects.

Table C.2: Comparison of cost items under imperfect forecasts.

Cost Item ($/month)	**Long-Term MPC** (Forecasted Load)	**Long-Term MPC** (Realized Load)	**Hierarchical MPC** (Imperfect Forecast)	**Standard MPC** (Perfect Forecast)	**Standard MPC** (Imperfect Forecast)
Total cost	126,690.05	135,334.06	135,374.74	136,230.18	138,411.67
Demand charge	128,594.86	137,220.12	137,220.12	138,057.22	140,330.69
Energy	-1,904.81	-1,886.06	-1,845.38	-1,827.04	-1,919.01

C.5 Conclusions

An approach to handle long horizons in MPC was proposed in the context of energy systems arising due to the need to capture long-term peak costs. In this approach, if periodicity constraints are enforced over short-term periods, the long horizon MPC problem can be posed as a stochastic programming problem with each period representing a scenario, the periodic state, and peak cost targets representing the design variables and the intra-stage operational policies representing the recourse variables. The SP setting reveals a mechanism to construct a hierarchical MPC scheme under which a long-term MPC planner provides state and peak cost targets to guide a short-term MPC controller. Through sim-

ulation case studies of a typical university campus with stationary battery storage, where the goal is to use the battery to decrease peak demand charges, it is shown that this hierarchical MPC scheme provides optimal operational policies under nominal (perfect forecast) conditions and can be extended to handle imperfect forecasts by correcting the short-term policies. It is also demonstrated through the simulation case study that the hierarchical MPC scheme yields improved performance over standard MPC schemes that use short-time horizons and long-term discounting factors. We have seen in Chapter 5 that the SP setting enables the use of advanced decomposition schemes (Higle and Sen, 1991), which can be used to progressively update the high-level MPC layer by adding cutting planes and enable a retroactive hierarchical scheme which only uses historical data to compute the optimal periodic states.

BIBLIOGRAPHY

Aguilera, R. P. and Quevedo, D. E. Stability analysis of quadratic MPC with a discrete input alphabet. *IEEE Transactions on Automatic Control*, 58(12):3190–3196, 2013.

Al-Mansour, F. and Kožuh, M. Risk analysis for CHP decision making within the conditions of an open electricity market. *Energy*, 32(10):1905–1916, 2007.

Alt, J., Anderson, M., and Jungst, R. Assessment of utility side cost savings from battery energy storage. *IEEE Transactions on Power Systems*, 12(3):1112–1120, 1997.

Appino, R. R., Ordiano, J. Á. G., Mikut, R., Faulwasser, T., and Hagenmeyer, V. On the use of probabilistic forecasts in scheduling of renewable energy sources coupled to storages. *Applied Energy*, 210:1207–1218, 2018.

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

Babu, C. and Ashok, S. Peak load management in electrolytic process industries. *IEEE Transactions on Power Systems*, 23(2):399–405, 2008.

Baccino, F., Conte, F., Massucco, S., Silvestro, F., and Grillo, S. Frequency regulation by management of building cooling systems through model predictive control. In *2014 Power Systems Computation Conference*, pages 1–7. IEEE, 2014.

Baldea, M. and Daoutidis, P. Control of integrated process networks – A multi-time scale perspective. *Computers & chemical engineering*, 31(5):426–444, 2007.

Baldea, M., Du, J., Park, J., and Harjunkoski, I. Integrated production scheduling and model predictive control of continuous processes. *AIChE Journal*, 61(12):4179–4190, 2015.

Bassett, M. H., Pekny, J. F., and Reklaitis, G. V. Decomposition techniques for the solution of large-scale scheduling problems. *AIChE Journal*, 42(12):3373–3387, 1996.

Beal, L. D., Clark, J. D., Anderson, M. K., Warnick, S., and Hedengren, J. D. Combined scheduling and control with diurnal constraints and costs using a discrete time formulation. *Proceedings of the FOCAPO (Foundations of Computer Aided Process Operations) and CPC (Chemical Process Control)*, pages 1–6, 2017.

Beal, L. D., Petersen, D., Grimsman, D., Warnick, S., and Hedengren, J. D. Integrated scheduling and control in discrete-time with dynamic parameters and constraints. *Computers & Chemical Engineering*, 115:361–376, 2018.

Bemporad, A. and Morari, M. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.

Bernardini, D. and Bemporad, A. Scenario-based model predictive control of stochastic constrained linear systems. In *Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference (CDC/CCC 2009)*, pages 6333–6338. IEEE, 2009.

Bertsekas, D. P. Incremental proximal methods for large scale convex optimization. *Mathematical Programming*, 129(2):163, 2011.

Bertsekas, D. P., Bertsekas, D. P., Bertsekas, D. P., and Bertsekas, D. P. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.

Bertsimas, D. and Tsitsiklis, J. N. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

Bezanson, J., Karpinski, S., Shah, V. B., and Edelman, A. Julia: A fast dynamic language for technical computing. *arXiv:1209.5145*, pages 1–27, 2012.

Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.

Birge, J. R. and Louveaux, F. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

Bıyık, E. and Arcak, M. Area aggregation and time-scale modeling for sparse nonlinear networks. *Systems & Control Letters*, 57(2):142–149, 2008.

Bonnans, J. F. and Shapiro, A. *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2013.

Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. *Time series analysis: Forecasting and control*. John Wiley & Sons, 2015.

Bradbury, K., Pratson, L., and Patino-Echeverri, D. Economic viability of energy storage systems based on price arbitrage potential in real-time U.S. electricity markets. *Applied Energy*, 114:512–519, 2014.

Braun, J. E. Reducing energy costs and peak electrical demand through optimal control of building thermal storage. *ASHRAE transactions*, 96(2):876–888, 1990.

Brunetto, C. and Tina, G. Optimal hydrogen storage sizing for wind power plants in day ahead electricity market. *IET Renewable Power Generation*, 1(4):220–226, 2007.

Camacho, E. F. and Alba, C. B. *Model predictive control*. Springer Science & Business Media, 2013.

Camacho, E. F., Ramírez, D. R., Limón, D., De La Peña, D. M., and Alamo, T. Model predictive control techniques for hybrid systems. *Annual Reviews in Control*, 34(1):21–31, 2010.

Cannon, M., Couchman, P., and Kouvaritakis, B. MPC for stochastic systems. *Lecture Notes in Control and Information Sciences*, 358(1):255–268, 2007.

Cao, Y., Laird, C. D., and Zavala, V. M. Clustering-based preconditioning for stochastic programs. *Computational Optimization and Applications*, 64(2):379–406, 2016.

CarøE, C. C. and Schultz, R. Dual decomposition in stochastic integer programming. *Operations Research Letters*, 24(1):37–45, 1999.

Chen, J. and Garcia, H. E. Economic optimization of operations for hybrid energy systems under variable markets. *Applied Energy*, 177:11–24, 2016.

Chiang, N.-Y. and Zavala, V. M. Large-scale optimal control of interconnected natural gas and electrical transmission systems. *Applied Energy*, 168:226–235, 2016.

Chicco, G. and Mancarella, P. Distributed multi-generation: A comprehensive view. *Renewable and Sustainable Energy Reviews*, 13:535–551, 2009.

Chow, J. H. and Kokotovic, P. V. Time scale modeling of sparse dynamic networks. *IEEE Transactions on Automatic Control*, 30(8):714–722, 1985.

Christidis, A., Koch, C., Pottel, L., and Tsatsaronis, G. The contribution of heat storage to the profitable operation of combined heat and power plants in liberalized electricity markets. *Energy*, 41(1):75–82, 2012.

Coderch, M., Willsky, A. S., Sastry, S. S., Castanon, D., and Others. Hierarchical aggregation of linear systems with multiple time scales. *IEEE Transactions on Automatic Control*, 28 (11):1017–1030, 1983.

Daoutidis, P., Lee, J. H., Harjunkoski, I., Skogestad, S., Baldea, M., and Georgakis, C. Integrating operations and control: A perspective and roadmap for future research. *Computers & Chemical Engineering*, 115:179–184, 2018.

de la Penad, D., Bemporad, A., and Alamo, T. Stochastic programming applied to model predictive control. In *44th IEEE Conference on Decision and Control, 2005 and 2005 European Control Conference (CDC-ECC 2005)*, pages 1361–1366. IEEE, 2005.

De Paepe, M. and Mertens, D. Combined heat and power in a liberalised energy market. *Energy Conversion and Management*, 48(9):2542–2555, 2007.

de Salis, R. T., Clarke, A., Wang, Z., Moyne, J., and Tilbury, D. Energy storage control for peak shaving in a single building. In *2014 IEEE PES General Meeting | Conference & Exposition*, pages 1–5. IEEE, 2014.

Delebecque, F., Quadrat, J. P., and Kokotovic, P. V. A unified view of aggregation and coherency in networks and Markov chains. *International Journal of Control*, 40(5):939–952, 1984.

Deng, K., Sun, Y., Li, S., Lu, Y., Brouwer, J., Mehta, P. G., Zhou, M., and Chakraborty, A. Model predictive control of central chiller plant with thermal energy storage via dynamic programming and mixed-integer linear programming. *IEEE Transactions on Automation Science and Engineering*, 12(2):565–579, 2014.

Di Cairano, S., Heemels, W. M. H., Lazar, M., and Bemporad, A. Stabilizing dynamic controllers for hybrid systems: a hybrid control lyapunov function approach. *IEEE Transactions on Automatic Control*, 59(10):2629–2643, 2014.

Dicorato, M., Forte, G., Pisani, M., and Trovato, M. Planning and operating combined wind-storage system in electricity market. *IEEE Transactions on Sustainable Energy*, 3(2): 209–217, 2012.

Diehl, M. and Bjornberg, J. Robust dynamic programming for min-max model predictive control of constrained uncertain systems. *IEEE Transactions on Automatic Control*, 49(12): 2253–2257, 2004.

DOE. Building energy data book 2011. *US Department of Energy*, 2011.

Domahidi, A., Ullmann, F., Morari, M., and Jones, C. N. Learning near-optimal decision rules for energy efficient building control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 7571–7576. IEEE, 2012.

Donadee, J. and Ilić, M. Stochastic co-optimization of charging and frequency regulation by electric vehicles. In *2012 North American Power Symposium (NAPS)*, pages 1–6. IEEE, 2012.

Donadee, J. and Ilić, M. D. Stochastic optimization of grid to vehicle frequency regulation capacity bids. *IEEE Transactions on Smart Grid*, 5(2):1061–1069, 2014.

Dong, X., Bao, G., Lu, Z., Yuan, Z., and Lu, C. Optimal battery energy storage system charge scheduling for peak shaving application considering battery lifetime. In *Informatics in Control, Automation and Robotics*, pages 211–218. Springer, 2011.

Dowling, A. W., Kumar, R., and Zavala, V. M. A multi-scale optimization framework for electricity market participation. *Applied Energy*, 190:147–164, 2017.

Drgoňa, J., Kvasnica, M., Klaučo, M., and Fikar, M. Explicit stochastic MPC approach to building temperature control. In *52nd IEEE Conference on Decision and Control*, pages 6440–6445. IEEE, 2013.

Dunning, I., Huchette, J., and Lubin, M. JuMP: A modeling language for mathematical optimization. *SIAM Review*, 59(2):295–320, 2017.

Dupačová, J., Consigli, G., and Wallace, S. W. Scenarios for multistage stochastic programs. *Annals of Operations Research*, 100(1-4):25–53, 2000.

Ekman, C. K. and Jensen, S. H. Prospects for large scale electricity storage in Denmark. *Energy Conversion and Management*, 51(6):1140–1147, 2010.

Fares, R. L., Meyers, J. P., and Webber, M. E. A dynamic model-based estimate of the value of a vanadium redox flow battery for frequency regulation in Texas. *Applied Energy*, 113: 189–198, 2014.

Farina, M., Giulioni, L., and Scattolini, R. Stochastic linear model predictive control with chance constraints – a review. *Journal of Process Control*, 44:53–67, 2016.

Feng, J., Brown, A., O'Brien, D., and Chmielewski, D. J. Smart grid coordination of a chemical processing plant. *Chemical Engineering Science*, pages 1–9, 2015.

Ferrarini, L., Mantovani, G., and Costanzo, G. T. A distributed model predictive control approach for the integration of flexible loads, storage and renewables. In *2014 IEEE 23rd International Symposium on Industrial Electronics (ISIE)*, pages 1700–1705. IEEE, 2014.

Foster, J. M. and Caramanis, M. C. Energy reserves and clearing in stochastic power markets: The case of plug-in-hybrid electric vehicle battery charging. In *49th IEEE Conference on Decision and Control (CDC)*, pages 1037–1044. IEEE, 2010.

García, C. E., Prett, D. M., and Morari, M. Model predictive control: Theory and practice – A survey. *Automatica*, 25(3):335–348, 1989.

Garifi, K., Baker, K., Touri, B., and Christensen, D. Stochastic model predictive control for demand response in a home energy management system. In *2018 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2018.

Geidl, M. and Andersson, G. Optimal power flow of multiple energy carriers. *IEEE Transactions on Power Systems*, 22(1):145–155, 2007.

Geoffrion, A. M. Generalized benders decomposition. *Journal of Optimization Theory and Applications*, 10(4):237–260, 1972.

González, J. L., Dimoulkas, I., and Amelin, M. Operation Planning of a CSP Plant in the Spanish Day-ahead Electricity Market. In *11th International Conference on the European Energy Market (EEM)*, pages 1–5, 2014.

Gray, A., Gao, Y., Lin, T., Hedrick, J. K., and Borrelli, F. Stochastic predictive control for semi-autonomous vehicles with an uncertain driver model. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pages 2329–2334. IEEE, 2013.

Grosso, J., Ocampo-Martínez, C., Puig, V., and Joseph, B. Chance-constrained model pre-

dictive control for drinking water networks. *Journal of Process Control*, 24(5):504–516, 2014.

Growe-Kuska, N., Heitsch, H., and Romisch, W. Scenario reduction and scenario tree construction for power management problems. In *2003 IEEE Bologna Power Tech Conference Proceedings,*, volume 3, pages 7–pp. IEEE, 2003.

Guigues, V. and Henrion, R. Joint dynamic probabilistic constraints with projected linear decision rules. *Optimization Methods and Software*, 32(5):1006–1032, 2017.

Gupta, D. and Maravelias, C. T. On deterministic online scheduling: Major considerations, paradoxes and remedies. *Computers & Chemical Engineering*, 94:312–330, 2016.

Gupta, D. and Maravelias, C. T. A general state-space formulation for online scheduling. *Processes*, 5(4):69, 2017.

Gupta, D., Maravelias, C. T., and Wassick, J. M. From rescheduling to online scheduling. *Chemical Engineering Research and Design*, 116:83–97, 2016.

Han, S., Han, S., and Sezaki, K. Development of an optimal vehicle-to-grid aggregator for frequency regulation. *IEEE Transactions on Smart Grid*, 1(1):65–72, 2010.

Hao, H., Middelkoop, T., Barooah, P., and Meyn, S. How demand response from commercial buildings will provide the regulation needs of the grid. In *50th Annual Allerton Conference on Communication, Control, and Computing*, pages 1908–1913, 2012.

Harjunkoski, I. and Grossmann, I. E. A decomposition approach for the scheduling of a steel plant production. *Computers & Chemical Engineering*, 25(11-12):1647–1660, 2001.

He, G., Chen, Q., Kang, C., Pinson, P., and Xia, Q. Optimal bidding strategy of battery storage in power markets considering performance-based regulation and battery cycle life. *IEEE Transactions on Smart Grid*, 7(5):2359–2367, 2015.

Helman, C. Thanks to fracking, natural gas supplies (barely) withstand 'polar vortex' assault. http://www.forbes.com/sites/christopherhelman/2014/01/08/

thanks-to-fracking-natural-gas-supplies-barely-withstand-polar-vortex-assault/#3679a50d5176, 2014.

Henze, G. P., Felsmann, C., and Knabe, G. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*, 43(2):173–183, 2004.

Henze, G. P., Biffar, B., Kohn, D., and Becker, M. P. Optimal design and operation of a thermal storage system for a chilled water plant serving pharmaceutical buildings. *Energy and Buildings*, 40(6):1004–1019, 2008.

Higle, J. L. and Sen, S. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research*, 16(3):650–669, 1991.

Hooshmand, A., Poursaeidi, M. H., Mohammadpour, J., Malki, H. A., and Grigoriads, K. Stochastic model predictive control method for microgrid management. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, pages 1–7, 2012.

Hu, T.-C., Rosalsky, A., and Volodin, A. On convergence properties of sums of dependent random variables under second moment and covariance restrictions. *Statistics & Probability Letters*, 78(14):1999–2005, 2008.

Huang, R., Harinath, E., and Biegler, L. T. Lyapunov stability of economically oriented NMPC for cyclic processes. *Journal of Process Control*, 21(4):501–509, 2011.

ISO New England. January 2014 FERC Data Request. Technical report, ISO New England, 2014. URL https://www.iso-ne.com/static-assets/documents/pubs/spcl_rpts/2014/iso_ne_response_ferc_data_request_january_2014.pdf.

Jackson, J. R. and Grossmann, I. E. Temporal decomposition scheme for nonlinear multisite production planning and distribution models. *Industrial & Engineering Chemistry Research*, 42(13):3045–3055, 2003.

Jaleeli, N., VanSlyck, L. S., Ewart, D. N., Fink, L. H., and Hoffmann, A. G. Understanding automatic generation control. *IEEE Transactions on Power Systems*, 7(3):1106–1122, 1992.

Jogwar, S. S., Baldea, M., and Daoutidis, P. Dynamics and control of process networks with large energy recycle. *Industrial & Engineering Chemistry Research*, 48(13):6087–6097, 2009.

Johnson, M. P., Bar-Noy, A., Liu, O., and Feng, Y. Energy peak shaving with local storage. *Sustainable Computing: Informatics and Systems*, 1(3):177–188, 2011.

Joshi, K. A. and Pindoriya, N. M. Day-ahead dispatch of battery energy storage system for peak load shaving and load leveling in low voltage unbalance distribution networks. In *2015 IEEE Power & Energy Society General Meeting*, pages 1–5. IEEE, 2015.

Kaut, M. and Wallace, S. W. *Evaluation of scenario-generation methods for stochastic programming*. Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II, Institut für Mathematik, 2003.

Kemp, J. U.S. power grid survived polar vortex, but only just. http://www.reuters.com/article/us-usa-power-weather-kemp-idUSKCN0HQ4TB20141003, 2014.

Kerrigan, E. C. *Robust constraint satisfaction: Invariant sets and predictive control*. PhD thesis, University of Cambridge, 2001.

Khalilpour, R. and Vassallo, A. Planning and operation scheduling of PV-battery systems: A novel methodology. *Renewable and Sustainable Energy Reviews*, 53:194–208, 2016.

Kim, K., Yang, F., Zavala, V. M., and Chien, A. A. Data centers as dispatchable loads to harness stranded power. *IEEE Transactions on Sustainable Energy*, 8(1):208–218, 2016.

Kim, K., Botterud, A., and Qiu, F. Temporal decomposition for improved unit commitment in power system production cost modeling. *IEEE Transactions on Power Systems*, 33(5): 5276–5287, 2018.

Kirby, B., O'Malley, M., Ma, O., Cappers, P., Corbus, D., Kiliccote, S., Onar, O., Starke, M., and Steinbergy, D. Load Participation in Ancillary Services (Workshop Report). Technical Report December, U.S. Deparment of Energy, 2011.

Kirches, C. *Fast numerical methods for mixed-integer nonlinear model-predictive control*. Springer, 2011.

Kobayashi, K., Shein, W. W., and Hiraishi, K. Large-scale MPC with continuous/discrete-valued inputs: Compensation of quantization errors, stabilization, and its application. *SICE Journal of Control, Measurement, and System Integration*, 7(3):152–158, 2014.

Kokotovic, P. V. Subsystems, time scales and multimodeling. *IFAC Proceedings Volumes*, 13 (6):xxvii–xxxiii, 1980.

Koller, R. W., Ricardez-Sandoval, L. A., and Biegler, L. T. Stochastic back-off algorithm for simultaneous design, control, and scheduling of multiproduct systems under uncertainty. *AIChE Journal*, 64(7):2379–2389, 2018.

Kost, C., Flath, C. M., and Möst, D. Concentrating solar power plant investment and operation decisions under different price and support mechanisms. *Energy Policy*, 61: 238–248, 2013.

Kottick, D., Blau, M., and Edelstein, D. Battery energy storage for frequency regulation in an island power system. *IEEE Transactions on Energy Conversion*, 8(3):455–459, 1993.

Kouvaritakis, B. and Cannon, M. *Stochastic Model Predictive Control*, pages 1–9. Springer London, London, 2013.

Kouzoupis, D., Klintberg, E., Diehl, M., and Gros, S. A dual newton strategy for scenario decomposition in robust multistage MPC. *International Journal of Robust and Nonlinear Control*, 28(6):2340–2355, 2018.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. A

stochastic model predictive control framework for stationary battery systems. *IEEE Transactions on Power Systems*, 33(4):4397, 2018a.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. A hierarchical model predictive control approach for battery systems. In *International High Performance Buildings Conference*, page Paper 260, 2018b.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. A stochastic dual dynamic programming framework for multiscale MPC. *IFAC-PapersOnLine*, 51(20):493–498, 2018c. 6th IFAC Conference on Nonlinear Model Predictive Control NMPC 2018.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. Handling long horizons in MPC: A stochastic programming approach. In *2018 Annual American Control Conference (ACC)*, pages 715–720. IEEE, 2018d.

Kumar, R., Jalving, J., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. Benchmarking stochastic and deterministic MPC: A case study in stationary battery systems. *AIChE Journal*, 65(7):e16551, 2019a.

Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., and Zavala, V. M. Hierarchical MPC schemes for periodic systems using stochastic programming. *Automatica*, 107:306–316, 2019b.

Kumar, R., Wenzel, M. J., ElBsat, M. N., Risbeck, M. J., Drees, K. H., and Zavala, V. M. Stochastic model predictive control for central HVAC plants. *Journal of Process Control*, 90:1–17, 2020.

Kwadzogah, R., Zhou, M., and Li, S. Model predictive control for HVAC systems – a review. In *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 442–447. IEEE, 2013.

Lara, C. L., Mallapragada, D. S., Papageorgiou, D. J., Venkatesh, A., and Grossmann, I. E. Deterministic electric power infrastructure planning: Mixed-integer programming model and nested decomposition algorithm. *European Journal of Operational Research*, 271 (3):1037–1054, 2018.

Lara, C. L., Siirola, J. D., and Grossmann, I. E. Electric power infrastructure planning under uncertainty: stochastic dual dynamic integer programming (SDDiP) and parallelization scheme. *Optimization and Engineering*, pages 1–39, 2019.

Leadbetter, J. and Swan, L. Battery storage system for residential electricity peak demand shaving. *Energy and Buildings*, 55:685–692, 2012.

Ledoit, O. and Wolf, M. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.

Lefort, A., Bourdais, R., Ansanay-Alex, G., and Guéguen, H. Hierarchical control method applied to energy management of a residential house. *Energy and Buildings*, 64:53–61, 2013.

Limon, D., Alvarado, I., Alamo, T., and Camacho, E. Robust tube-based MPC for tracking of constrained linear systems with additive disturbances. *Journal of Process Control*, 20 (3):248–260, 2010.

Lin, J., Leung, K. C., and Li, V. O. K. Optimal Scheduling with Vehicle-to-Grid Regulation Service. *IEEE Internet of Things Journal*, 1(6):556–569, 2014.

Linderoth, J., Shapiro, A., and Wright, S. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, 142(1):215–241, 2006.

Lizarraga-Garcia, E., Ghobeity, A., Totten, M., and Mitsos, A. Optimal operation of a solar-thermal power plant with energy storage and electricity buy-back from grid. *Energy*, 51: 61–70, 2013.

Lopez-Negrete, R., DÕAmato, F. J., Biegler, L. T., and Kumar, A. Fast nonlinear model predictive control: Formulation and industrial process applications. *Computers & Chemical Engineering*, 51:55–64, 2013.

Lorenzen, M., Dabbene, F., Tempo, R., and Allgöwer, F. Constraint-tightening and stability in stochastic model predictive control. *IEEE Transactions on Automatic Control*, 62(7): 3165–3177, 2016.

Lu, C., Xu, H., Pan, X., and Song, J. Optimal sizing and control of battery energy storage system for peak load shaving. *Energies*, 7(12):8396–8410, 2014.

Lucas, A. and Chondrogiannis, S. Smart grid energy storage controller for frequency regulation and peak shaving, using a vanadium redox flow battery. *International Journal of Electrical Power and Energy Systems*, 80:26–36, 2016.

Lucia, S., Finkler, T., and Engell, S. Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty. *Journal of Process Control*, 23 (9):1306–1319, 2013.

Lucia, S., Andersson, J. A., Brandt, H., Diehl, M., and Engell, S. Handling uncertainty in economic nonlinear model predictive control: A comparative case study. *Journal of Process Control*, 24(8):1247–1259, 2014.

Ma, J., Qin, J., Salsbury, T., and Xu, P. Demand reduction in building energy systems based on economic model predictive control. *Chemical Engineering Science*, 67(1):92–100, 2012a.

Ma, Y. *Model predictive control for energy efficient buildings*. PhD thesis, UC Berkeley, 2012.

Ma, Y. and Borrelli, F. Fast stochastic predictive control for building temperature regulation. In *2012 American Control Conference (ACC)*, pages 3075–3080. IEEE, 2012.

Ma, Y., Borrelli, F., Hencey, B., Coffey, B., Bengea, S., and Haves, P. Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology*, 20(3):796–803, 2012b.

Ma, Y., Matuško, J., and Borrelli, F. Stochastic model predictive control for building HVAC systems: Complexity and conservatism. *IEEE Transactions on Control Systems Technology*, 23(1):101–116, 2014.

Madaeni, S. H., Sioshansi, R., and Denholm, P. How thermal energy storage enhances the economic viability of concentrating solar power. *Proceedings of the IEEE*, 100(2):335–347, 2011.

Magni, L. and Scattolini, R. Robustness and robust design of MPC for nonlinear discrete-time systems. In *Assessment and Future Directions of Nonlinear Model Predictive Control*, pages 239–254. Springer, 2007.

Mayne, D. Q., Rawlings, J. B., Rao, C. V., and Scokaert, P. O. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

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

Mercier, P., Cherkaoui, R., and Oudalov, A. Optimizing a battery energy storage system for frequency control application in an isolated power system. *IEEE Transactions on Power Systems*, 24(3):1469–1477, 2009.

Mesbah, A. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.

Mesbah, A., Streif, S., Findeisen, R., and Braatz, R. D. Stochastic nonlinear model predictive control with probabilistic constraints. In *2014 American control conference*, pages 2413–2419. IEEE, 2014.

Mitra, S., Sun, L., and Grossmann, I. E. Optimal scheduling of industrial combined heat and power plants under time-sensitive electricity prices. *Energy*, 54:194–211, 2013.

Mohsenian-Rad, H. Optimal bidding, scheduling, and deployment of battery systems in California day-ahead energy market. *IEEE Transactions on Power Systems*, 31(1):442–453, 2016.

Moura, S. J., Callaway, D. S., Fathy, H. K., and Stein, J. L. Tradeoffs between battery energy capacity and stochastic optimal power management in plug-in hybrid electric vehicles. *Journal of Power Sources*, 195(9):2979–2988, 2010.

Nghiem, T. X. and Jones, C. N. Data-driven demand response modeling and control of buildings with gaussian processes. In *2017 American Control Conference (ACC)*, pages 2919–2924. IEEE, 2017.

Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.

Oldewurtel, F., Jones, C. N., Parisio, A., and Morari, M. Stochastic model predictive control for building climate control. *IEEE Transactions on Control Systems Technology*, 22(3):1198–1205, 2013.

Oudalov, A., Chartouni, D., and Ohler, C. Optimizing a battery energy storage system for primary frequency control. *IEEE Transactions on Power Systems*, 22(3):1259–1266, 2007a.

Oudalov, A., Cherkaoui, R., and Beguin, A. Sizing and optimal operation of battery energy storage system for peak shaving application. In *2007 IEEE Lausanne POWERTECH, Proceedings*, pages 621–625, 2007b.

Oudalov, A., Chartouni, D., Ohler, C., and Linhofer, G. Value analysis of battery energy storage applications in power systems. In *2006 IEEE PES Power Systems Conference and Exposition*, pages 2206–2211. IEEE, 2006.

Pappas, G. J., Lafferriere, G., and Sastry, S. Hierarchically consistent control systems. *IEEE Transactions on Automatic Control*, 45(6):1144–1160, 2000.

Patel, N. R. and Rawlings, J. B. Applications of MPC to building HVAC systems. In *Handbook of Model Predictive Control*, pages 607–623. Springer, 2019.

Patel, N. R., Risbeck, M. J., Rawlings, J. B., Wenzel, M. J., and Turney, R. D. Distributed economic model predictive control for large-scale building temperature regulation. In *American Control Conference (ACC), 2016*, pages 895–900. IEEE, 2016.

Pattison, R. C., Touretzky, C. R., Harjunkoski, I., and Baldea, M. Moving horizon closed-loop production scheduling using dynamic process models. *AIChE Journal*, 63(2):639–651, 2017.

Paulson, J. A., Streif, S., and Mesbah, A. Stability for receding-horizon stochastic model predictive control. In *2015 American Control Conference (ACC)*, pages 937–943. IEEE, 2015.

Paulson, J. A., Buehler, E. A., Braatz, R. D., and Mesbah, A. Stochastic model predictive control with joint chance constraints. *International Journal of Control*, pages 1–14, 2017.

Pereira, M. V. and Pinto, L. M. Multi-stage stochastic optimization applied to energy planning. *Mathematical programming*, 52(1-3):359–375, 1991.

Pereira, M. Optimal stochastic operations scheduling of large hydroelectric systems. *International Journal of Electrical Power & Energy Systems*, 11(3):161–169, 1989.

Philpott, A. B. and Guan, Z. On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters*, 36(4):450–455, 2008.

Picasso, B., De Vito, D., Scattolini, R., and Colaneri, P. An MPC approach to the design of two-layer hierarchical control systems. *Automatica*, 46(5):823–831, 2010.

Qin, S. J. and Badgwell, T. A. An overview of nonlinear model predictive control applications. In *Nonlinear Model Predictive Control*, pages 369–392. Springer, 2000.

Qin, S. J. and Badgwell, T. A. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764, 2003.

Ragan, P. and Manuel, L. Statistical extrapolation methods for estimating wind turbine extreme loads. *Journal of Solar Energy Engineering*, 130(3):031011, 2008.

Rahimi, A., Zarghami, M., Vaziri, M., and Vadhva, S. A simple and effective approach for peak load shaving using battery storage systems. In *2013 North American Power Symposium (NAPS)*, pages 1–5. IEEE, 2013.

Rao, C. V., Rawlings, J. B., and Mayne, D. Q. Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations. *IEEE Transactions on Automatic Control*, 48(2):246–258, 2003.

Rastler, D. *Electricity energy storage technology options: a white paper primer on applications, costs and benefits*. Electric Power Research Institute, 2010.

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

Rawlings, J. B., Patel, N. R., Risbeck, M. J., Maravelias, C. T., Wenzel, M. J., and Turney, R. D. Economic MPC and real-time decision making with application to large-scale HVAC energy systems. *Computers & Chemical Engineering*, 114:89–98, 2018.

Rawlings, J. B., Mayne, D. Q., and Diehl, M. *Model predictive control: Theory, computation, and design*, volume 2. Nob Hill Publishing, 2017.

Renteria, J. A., Cao, Y., Dowling, A. W., and Zavala, V. M. Optimal PID controller tuning using stochastic programming techniques. *AIChE Journal*, 64(8):2997–3010, 2018.

Risbeck, M. J., Maravelias, C. T., Rawlings, J. B., and Turney, R. D. Cost optimization of combined building heating/cooling equipment via mixed-integer linear programming. In *American Control Conference (ACC), 2015*, pages 1689–1694. IEEE, 2015.

Risbeck, M. J., Maravelias, C. T., Rawlings, J. B., and Turney, R. D. Closed-loop scheduling for cost minimization in HVAC central plants. In *International High Performance Buildings Conference*, page Paper 177, 2016.

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

Römisch, W. Scenario reduction techniques in stochastic programming. In *International Symposium on Stochastic Algorithms*, pages 1–14. Springer, 2009.

Rong, A. and Lahdelma, R. Efficient algorithms for combined heat and power production planning under the deregulated electricity market. *European Journal of Operational Research*, 176(2):1219–1245, 2007.

Rudin, W. *Principles of mathematical analysis*. McGraw-Hill New York, 3rd edition, 1976.

Sarker, M. R., Dvorkin, Y., and Ortega-Vazquez, M. A. Optimal Participation of an Electric Vehicle Aggregator in Day-Ahead Energy and Reserve Markets. *IEEE Transactions on Power Systems*, PP(99):1–10, 2015.

Scattolini, R. Architectures for distributed and hierarchical model predictive control – A review. *Journal of Process Control*, 19(5):723–731, 2009.

Scattolini, R. and Colaneri, P. Hierarchical model predictive control. In *2007 46th IEEE Conference on Decision and Control*, pages 4803–4808. IEEE, 2007.

Sebastián, R. Application of a battery energy storage for frequency regulation and peak shaving in a wind diesel power system. *IET Generation, Transmission & Distribution*, 10 (3):764–770, 2016.

Shahidehpour, M., Fu, Y., and Wiedman, T. Impact of natural gas infrastructure on electric power systems. *Proceedings of the IEEE*, 93(5):1042–1056, 2005.

Shapiro, A. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, 2011.

Shapiro, A. and Wardi, Y. Convergence analysis of gradient descent stochastic algorithms. *Journal of Optimization Theory and Applications*, 91(2):439–454, 1996.

Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: Modeling and theory*. SIAM, 2009.

Shi, Y., Xu, B., Wang, D., and Zhang, B. Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains. *IEEE Transactions on Power Systems*, 33(3):2882–2894, 2017.

Shin, S. and Zavala, V. M. Multi-grid schemes for multi-scale coordination of energy systems. In *Energy Markets and Responsive Grids*, pages 195–222. Springer, 2018.

Sifuentes, W. S. and Vargas, A. Hydrothermal scheduling using benders decomposition: Accelerating techniques. *IEEE Transactions on Power Systems*, 22(3):1351–1359, 2007.

Sigrist, L., Lobato, E., and Rouco, L. Energy storage systems providing primary reserve and peak shaving in small isolated power systems: An economic assessment. *International Journal of Electrical Power & Energy Systems*, 53:675–683, 2013.

Simon, H. A. and Ando, A. Aggregation of variables in dynamic systems. *Econometrica: Journal of the Econometric Society*, pages 111–138, 1961.

Singh, S., Singh, S. K., Chanana, S., and Singh, Y. Frequency regulation of an isolated hybrid power system with battery energy storage system. In *2014 Power and Energy Systems: Towards Sustainable Energy*, pages 1–6. IEEE, 2014.

Sioshansi, R., Denholm, P., Jenkin, T., and Weiss, J. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics*, 31(2):269–277, 2009.

Sortomme, E. and El-Sharkawi, M. A. Optimal scheduling of vehicle-to-grid energy and ancillary services. *IEEE Transactions on Smart Grid*, 3(1):351–359, 2012.

Subramanian, K., Maravelias, C. T., and Rawlings, J. B. A state-space model for chemical production scheduling. *Computers & Chemical Engineering*, 47:97–110, 2012.

Subramanian, K., Rawlings, J. B., and Maravelias, C. T. Economic model predictive control for inventory management in supply chains. *Computers & Chemical Engineering*, 64:71–80, 2014.

The White House. Office of the Press Secretary: President Obama signs executive order promoting industrial energy efficiency. https://obamawhitehouse.archives.gov/the-press-office/2012/08/30/president-obama-signs-executive-order-promoting-industrial-energy-effici, 2012.

Usaola, J. Operation of concentrating solar power plants with storage in spot electricity markets. *IET Renewable Power Generation*, 6(1):59–66, 2012.

Vagropoulos, S. I. and Bakirtzis, A. G. Optimal bidding strategy for electric vehicle aggregators in electricity markets. *IEEE Transactions on Power Systems*, 28(4):4031–4041, 2013.

Walawalkar, R., Apt, J., and Mancini, R. Economics of electric energy storage for energy arbitrage and regulation in New York. *Energy Policy*, 35(4):2558–2568, 2007.

Walling, R. A., Banunarayanan, V., Chahal, A., Freeman, L., Martinez, J., Miller, N., Van Zandt, D., Walling, M., and Walling, R. Analysis of wind generation impact on ercot ancillary services requirements. *GE Energy, GE Project Team*, 2008.

Wang, Y., Wang, B., Chu, C. C., Pota, H., and Gadh, R. Energy management for a commercial building microgrid with stationary and mobile battery storage. *Energy and Buildings*, 116:141–150, 2016.

White, C. D. and Zhang, K. M. Using vehicle-to-grid technology for frequency regulation and peak-load reduction. *Journal of Power Sources*, 196(8):3972–3980, 2011.

Zakeri, B. and Syri, S. Economy of electricity storage in the Nordic electricity market: The case for Finland. In *11th International Conference on the European Energy Market (EEM14)*, pages 1–6, 2014.

Zanon, M., Grüne, L., and Diehl, M. Periodic optimal control, dissipativity and MPC. *IEEE Transactions on Automatic Control*, 62(6):2943–2949, 2017.

Zavala, V. M. Stochastic optimal control model for natural gas networks. *Computers & Chemical Engineering*, 64:103–113, 2014.

Zavala, V. M. New Architectures for Hierarchical Predictive Control. *IFAC-PapersOnLine*, 49(7):43–48, 2016. 11th IFAC Symposium on Dynamics and Control of Process SystemsIncluding Biosystems DYCOPS-CAB 2016.

Zavala, V. M., Laird, C. D., and Biegler, L. T. Interior-point decomposition approaches for parallel solution of large-scale nonlinear parameter estimation problems. *Chemical Engineering Science*, 63(19):4834–4845, 2008.

Zavala, V. M., Anitescu, M., and Krause, T. On the optimal on-line management of photovoltaic-hydrogen hybrid energy systems. In *Proceedings of 10th International Symposium on Process Systems Engineering*, pages 2579–2586, 2009a.

Zavala, V. M., Constantinescu, E. M., Krause, T., and Anitescu, M. On-line economic optimization of energy systems using weather forecast information. *Journal of Process Control*, 19(10):1725–1736, 2009b.

Zhang, X., Schildbach, G., Sturzenegger, D., and Morari, M. Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty. In *2013 European Control Conference (ECC)*, pages 1029–1034. IEEE, 2013.

Zhao, P., Henze, G. P., Plamp, S., and Cushing, V. J. Evaluation of commercial building HVAC systems as frequency regulation providers. *Energy and Buildings*, 67:225–235, 2013.

Zhao, P., Henze, G. P., Brandemuehl, M. J., Cushing, V. J., and Plamp, S. Dynamic frequency regulation resources of commercial buildings through combined building sys-

tem resources using a supervisory control methodology. *Energy and Buildings*, 86:137–150, 2015.

Zou, J., Ahmed, S., and Sun, X. A. Nested decomposition of multistage stochastic integer programs with binary state variables. *Optimization Online*, 5436, 2016.

Zou, J., Ahmed, S., and Sun, X. A. Stochastic dual dynamic integer programming. *Mathematical Programming*, 175(1-2):461–502, 2019.