# Modelling, Design, and Control of Energy Systems: A Data-Driven Approach

by

Fiodar Kazhamiaka

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2019

**Examining Committee Membership**

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

| | |
|---|---|
| Supervisor(s): | Srinivasan Keshav<br>Professor, Cheriton School of Computer Science,<br>University of Waterloo |
| | Catherine Rosenberg<br>Professor, Dept. of Electrical and Computer Engineering,<br>University of Waterloo |
| Internal Member: | Bernard Wong<br>Professor, Cheriton School of Computer Science,<br>University of Waterloo |
| Internal Member: | Yaoliang Yu<br>Professor, Cheriton School of Computer Science,<br>University of Waterloo |
| Internal-External Member: | Claudio Cañizares<br>Professor, Dept. of Electrical and Computer Engineering,<br>University of Waterloo |
| External Member(s): | Steven Low<br>Professor, Dept. of Computing and Mathematical Sciences,<br>California Institute of Technology |

## Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

## Abstract

In 2018, nearly two-thirds of newly installed global power generation has come from renewable energy sources. Distributed installations of solar photovoltaic (PV) panels have been at the forefront of this global energy transition. In many places, the cost of solar power has dropped below the cost of fossil fuels such as coal. The main challenge in incorporating this growing source of clean and cheap energy is its high variability; it must often be used in conjunction with an expensive energy storage system to help match electricity supply and demand. Despite the growing focus on energy storage and its role in helping meet the ambitious renewable energy targets set by climate-conscious policy makers, the relatively high capital cost of combined PV-storage systems has limited their widespread adoption.

The high cost of PV-storage systems may be offset by the value they provide to system owners. The combination of PV panel and energy storage components adds complexity and flexibility to an energy system where both supply and demand are stochastic and depend on many factors, and this contributes to the technical challenges of system design and operation. Practical methods for increasing the value of PV-storage systems through effective system design and operation are the focus of this dissertation. The proposed approach to solving these problems involves the use of system models and data describing the system's operating environment. Our research consists of two main components: the theoretical component, i.e., modelling, and the practical component, i.e., analysis of energy systems on the basis of models and data.

In this thesis, we construct new battery models that enable simulation and optimization studies of PV-storage systems. These models are then used to develop several advanced methods for designing and operating PV-storage systems based on available solar generation and electricity consumption data. We study the problem of determining the combined sizes of PV panel and storage components to meet a given system load target at the lowest possible cost. We also study the problem of system operation, with the objective of increasing system value via minimization of operating expenses. The sizing and control methods developed in this thesis are based primarily on system simulation, mathematical programming, and neural networks, and are evaluated on datasets of PV generation and electricity consumption measurements of buildings.

Among our contributions are an accurate battery model for simulation studies that can be easily calibrated. We further derive models for optimization studies with various degrees of accuracy and complexity, including a linear model with higher accuracy than existing linear models. For system sizing, we develop a novel approach for sizing a PV-storage system to reliably meet a load performance target at the lowest possible cost. For

system operation, we develop a set of algorithms which achieve high performance with low information requirements, a mixed-timescale approach to reduce the online computational complexity of model predictive control, and a system controller designed to encode a deep neural network with a model predictive control policy and capable of refining its performance over time while adapting to changes in the operating environment.

**Acknowledgements**

Firstly, I thank my advisors Catherine Rosenberg and Srinivasan Keshav for creating a healthy and engaging environment, for inspiring me to work diligently, and for their endless belief and support. I thank Catherine for relentlessly demonstrating the importance of asking tough questions. I thank Keshav for helping me develop the courage to follow my instincts. This dissertation would not have been possible without their guidance.

I thank my committee members Steven Low, Claudio Canizares, Bernard Wong, and Yaoliang Yu, whose comments and advice have helped me develop a better understanding of my work.

I thank Khuzaima Daudjee and Bernard Wong (again) for allowing me the opportunity to broaden the scope of my expertise and dabble in networking and distributed systems research. Our conversations have been valuable in shaping my understanding of what it means to work in academia.

I thank my colleagues in Germany, namely Karl-Heinz Pettinger, Hermann de Meer, Patrick Jochem, Robert Basmadjian, and Florian Niedermeier, who have been very welcoming, enthusiastic, and productive in our work together.

I thank my colleagues at the ISS4E research group for the lively and engaging discussions on research, life, and everything in between. In particular, I wish to thank Costin Ograda-Bratu, Omid Ardakanian, Rayman Preet, Tommy Carpenter, Reid Miller, Adedamola Adepetu, Ankit Pat, Alimohammad Rabbani, Christian Gorenflo, Michael Doroshenko, Come Carquex, Yerbol Aussat, Ansis Rosmanis, Sun Sun, Dimcho Karakashev, and Kayla Hardie. I especially thank Yashar Ghiassi-Farrokhfal, who has been an excellent mentor since my undergraduate research days. I am lucky to have gone through graduate school among these fine people.

I thank my family for their encouragement and assistance in completing my graduate studies. I thank my mother for providing nutritional support, my brother and sister for their technical consultations, and my father for his writing advice. I especially thank my wife for going the extra mile, for encouraging me all the way through graduate school, and for the sacrifices she has made to support me along the way.

Finally, I thank the Natural Sciences and Engineering Council of Canada, as well as the Cheriton School of Computer Science, for supporting me with funding for this research.

## Dedication

This dissertation is dedicated to my brother Mikhail and sister Hanna, who may yet decide to write their own.

# Table of Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

Energy storage has played a role in electrical power systems for over a century. The main purpose of storing energy is to decouple the time of electricity supply from the time of demand, introducing flexibility into both generation and consumption processes. Recently, there has been a growing focus on energy storage and its role in meeting the renewable energy targets set by climate-conscious policymakers at various levels of government. Energy storage is considered to be a crucial element for increasing the penetration of renewable generation sources, in particular, solar photovoltaic (PV) panels and wind turbines. A high penetration of renewable electricity sources is one of the key aspects of a low-carbon electricity grid, and the intermittency of these sources is the primary barrier to their widespread adoption.

The use of energy storage systems to aid in maintaining the balance of supply and demand in the electricity grid has so far been on a very limited scale; the total energy storage capacity expressed as a fraction of daily average load is mere 0.2% in the United States [11, 5] and even lower in Canada [176, 39]. The majority of this storage is in the form of large pumped-hydro storage facilities [11], which pump water to an elevated reservoir when electricity supply exceeds demand and use its gravitational potential energy to spin turbines when demand exceeds supply. The increased deployment of pumped-hydro storage is limited by its high capital cost and restricting geographical requirements.

During the last decade, there has been a growing trend in the deployment of electrochemical batteries [6], especially Lithium-ion technologies, resulting in an increase in energy storage capacity. The early adopters of battery systems have been power producing-companies who purchase large-scale systems to directly support the electrical grid. However, storage deployment in the *end-user environment* is growing as well. Batteries are

being installed behind-the-meter[1] in residential, commercial, and industrial buildings.

A battery adds flexibility to a building's energy system, and this flexibility can be utilized in several ways. For example, in conjunction with a local PV installation, the battery may store excess PV energy and use it at night, when PV power is unavailable. Besides, the battery may be used to shift the system's net grid load to times where grid electricity is cheaper, which often correlate with periods during which the electricity is generated with low-cost methods. Other, more advanced applications of distributed energy generation and storage systems are also being explored, such as taking part in grid demand-response markets by aggregating many buildings' batteries into a virtual power plant [123], or selling energy to a neighbor through a peer-to-peer energy sharing network [105, 28].

The analysis of these trends[2] naturally leads to the notion of a PV-storage system, in which PV panels and batteries are integrated. Such a system creates opportunities to significantly offset the owner's reliance on energy from the grid, and increase the amount of clean energy consumption. While this consideration may convince certain environmentally-conscious end-users to purchase PV-storage systems, it is not sufficient for the majority of potential consumers who base their purchasing decisions primarily on the comparison between system cost and payback [4], or, in other terms, on the expected return-on-investment (ROI).

The cost of PV-storage systems is an important factor in their adoption. The costs of both PV panels and batteries have been in steady decline [46, 124]. PV panel prices vary across different parts of the world; at the time this research was conducted, residential installations typically have costs in the range of 1000-3000 USD per kWp. Batteries have a price range of 500-1000 USD per kWh of installed storage capacity. Hence, a system that is capable of producing a noticeable effect on the energy consumption of a typical household requires a substantial investment. This is the limiting factor in the widespread adoption of PV-storage systems.

The goal of this dissertation is to increase the ROI of end-user PV-storage systems and thereby increase their adoption. There are many different approaches to address this challenge. A chemist or materials engineer may seek to develop new and improved battery technologies or photovoltaic materials. An economist may study the effectiveness of various incentive schemes that encourage individuals to invest in these systems. The research presented in this thesis lies at the intersection of computer science, mathematics, and electrical engineering, and utilizes the typical methodologies derived from these disciplines.

[1] Behind-the-meter refers to the point of electricity consumption. The electricity meter used to bill the user for their consumption sees the net effect of user consumption and system operation.

[2] Referring to the trends in distributed PV panel and distributed battery installations.

There are many different stakeholders interested in PV-storage system adoption, who may each have distinct perspectives and goals. For instance, governments are interested in meeting their carbon reduction targets. Electrical utility companies and grid operators are interested in using PV generation and storage to reduce their costs. System manufacturers are interested in selling systems to make a profit. System owners are interested in reducing their electricity bills. In this dissertation, we take the perspective of PV-storage system owners who wish to satisfy their electricity demand at the lowest possible cost. In other words, this work focuses on increasing PV-storage system value for end-consumers.

We consider two main approaches to the increase of value, and thereby the ROI, of PV-storage systems.

1. **Choosing the right sizes for system components**. This involves sizing the components of the system, specifically the battery and PV panels. The goal is to find the optimal system size, i.e., the one that achieves a desired system performance target at the lowest possible cost.

2. **Operating the system in the best possible way**. Energy storage provides flexibility in energy consumption and creates opportunities for optimization. The goal is to minimize the cost of operation; for end-user systems, this means minimizing the payments to the utility company.

To optimize the sizing and operation of PV-storage systems, we integrate data traces consisting of localized PV generation and load measurements. Together, such datasets describe a PV-storage system's operating environment. The combination of this data with a dynamic system model enables the study of the system over potential state trajectories, as described by the data. We use this approach to develop methods for solving system design and control problems. Undertaking this work requires us to address some fundamental challenges, as described next.

## 1.1 Challenges

We identify four fundamentally challenging aspects of undertaking a data-driven study of PV-storage system optimization.

1. **System Complexity**: Studying the optimization of PV-storage system size and operation requires a battery model. The model should be accurate, easy to calibrate

and be usable as part of an optimization study without incurring overwhelming computational expenses. Thus, the development of battery models is a fundamental part of our work. Lithium-ion batteries are complex electrochemical systems. Accurately modelling the relationship between aspects such as battery voltage, energy content, applied current, energy conversion efficiency, and others, is a challenging task. Creating models which are tractable enough to be used as part of a system optimization study, without sacrificing heavily on model accuracy, is another major challenge.

2. **Coupling of Design and Control**: The optimal sizing of a PV-storage system depends on how it is operated, and optimal operation depends on the sizing. This coupling makes it difficult to simultaneously design both an optimal sizing and optimal operation strategy. Other aspects are also coupled, for example in the choice of specific PV panel and battery technologies, and the choice of the applications that the system is used for. Studying this problem jointly is a very difficult task due to the complexity of the system and the number of inter-dependent system variables. We discuss our approach to breaking the coupling in Section 2.2.

3. **Unknown Behaviour**: Configuring a PV-storage system with an appropriate sizing and operating strategy is typically done before the system is deployed. Once deployed, the system operates in an environment with many unknowns. This environment is highly dynamic, in part due to the unpredictability of the people who use the system to meet their energy demands, and in part due to the many changes that electrical systems are currently undergoing. For these reasons, ensuring that the system size and operating strategy is robust and adaptable to the future environment is challenging.

4. **Limited Data**: When a PV-storage system is deployed, there is often very little data available to support the configuration of its sizing and operating strategy. While we have a generally good idea of how much PV generation to expect from a particular location, the specific patterns of the system owner's electrical load are often unclear. In recent years, the deployment of smart meters has allowed some homeowners to get access to this data. However, given the stochasticity and non-stationarity of both solar and load processes (discussed in Section 4.2.2), several years of data are often needed to accurately represent these processes. Considering the lack of historical data combined with the fact that the future is unknown, it is difficult to provide guarantees on PV-storage system performance.

To address these challenges, we make use of a wide variety of methods. Our work features system simulation, mathematical programming, probability bounds, rule-based

control algorithms, model predictive control, neural networks, experiments on real batteries, and an evaluation on real measurement data.

## 1.2 Dissertation Outline

The topic of this research is related to the critical, complex, and multifaceted problem of anthropogenic climate change. Our work is relevant to the sub-problem of renewable energy adoption and integration, and we hope that it is a step forward to addressing the related problem of PV-storage system ROI. Where appropriate, the descriptions of our specific contributions are accompanied by a series of propositions regarding potential further studies.

The layout of this dissertation is as follows. Definitions of system terminology and overview of the research paradigm are provided in Chapter 2. This is followed by the three central parts of the thesis. In Chapter 3, we establish a theoretical understanding of the system through the development of a mathematical model. In Chapters 4 and 5, we bring the model into practice through the development of a framework for the optimization of system design and for the optimization of system operation, respectively. A brief description of each of these central parts follows.

### 1.2.1 Mathematical Models for Storage Systems

A mathematical model is a crucial element for the study of system optimization. A data-driven study of PV-storage systems, where power generation and consumption is modelled via data traces, requires a battery model. Among the many desirable properties of such models, the focus of our work is on the ease of model calibration, and on the trade-off between accuracy and complexity. The latter refers to the careful pruning of complex battery model features to make it tractable while minimizing losses to modelling accuracy. In Chapter 3, we develop a set of models which can be calibrated using only the data typically available in a battery manufacturer's specifications sheet, and which explore different trade-offs between accuracy and complexity.

Among our contributions is an accurate and easily calibrated model which is suitable for battery system simulation, and a family of analytic battery models for energy system optimization studies, including a linear model with higher accuracy than existing linear models. The simulation model has been implemented as a Matlab system block that is

compatible with Simulink simulation software. Our work on battery models has been published as [81, 78, 80]

### 1.2.2 PV-Storage System Sizing

Someone purchasing a PV-storage system for partially or completely offsetting grid usage is faced with the following question: how much of each should they buy? The task is to choose the combination of PV panels and battery modules from a discrete set of available sizes. It is desirable to purchase the cheapest system that meets the user's needs in order to maximize ROI. In Chapter 4, we develop a novel framework for computing a robust system sizing based on data-driven system simulation and probability bounds. In contrast with other existing work, we relax the assumption that the available historical data which characterizes the system's operating environment is perfectly representative of the future. Instead, we use a weaker assumption that the dynamic properties of both the past and future operating environments are characterized by the same probability distribution.

### 1.2.3 Optimization of System Operation

To maximize PV-storage system ROI, the system must be operated to minimize the cost of providing power to the user. This can be done in many ways, for example by charging and discharging the battery to take advantage of differences in the price of grid electricity throughout the day. The effectiveness of an operating strategy depends heavily on the operating environment, including the grid pricing scheme, load and PV generation patterns, as well as on system size. In Chapter 5, we develop control strategies for PV-storage systems under two grid pricing schemes: time-of-use pricing and peak-demand pricing. Among our contributions is (a) a mixed-timescale model predictive control approach that achieves quasi-optimal performance with relatively low computational hardware requirements, and (b) a deep neural network controller that is capable of learning an effective control policy over time and adapting to changes in the operating environment. Our work on system operation has been published as [79, 77].

This high-level description of our methods and objectives leads us to the description of the methodology used in this dissertation.

## 1.3 Research Methodology

The methodological considerations of this work include the recognition of the complexity of data, the plurality of methods, and the practicality of results. The primary technique of our research is the data-driven analysis of mathematical system models. Accordingly, the theoretical component of this work consists in the formulation of models. In the context of this work, we use the word "model" to refer to a simplified mathematical description of a system or process, meant to assist calculations and predictions [149]. Once the models are developed, we use them as a framework for the analysis of energy systems to obtain practical conclusions or recommendations. Thus, this research consists of two main components: the theoretical component, i.e., modelling, and the practical component, i.e., analysis of energy systems based on models and data.

Many of the modern methods for energy system analysis are driven by data: measurements of dynamic system parameters over time. The time series corresponding to the power output of a PV panel or the instantaneous power consumption of a building are examples of such data. This data can be used to develop realistic system models, which may then be analyzed to draw conclusions about the system under study.

We use real measurement data wherever possible to model relevant system processes. An alternative is to use data generation techniques based on mathematical models of the process, which are often extracted from data. These two types of data sources are referred to as data-driven and model-driven, respectively. The real data reflect the entire variety of factors affecting the process, including possible correlation with other processes[3]; in place of real data, a very complex model is required to mimic the processes accurately. In this research, the preference for the data-driven approach is based on the supposition that real data more accurately reflects the reality in comparison with a mathematical model.

The fundamental assumption of many data-driven studies is that historical data is representative of the future. Without this assumption, data would be useless as a base for *a posteriori* analysis. Indeed, historical data contain hints about future trends in the corresponding processes. However, besides the valuable information which may be obtained through data analysis, two other possibilities must be taken into account upon working with data. First, the data may contain information which is not representative of the future. Examples of this could be *noise* which distorts the true signal that the data is meant to convey, and *outliers* which skew the results of analysis toward the extreme rather than moderate or expected cases. Second, the data may be missing information which is

[3] An example of correlation is the increase in power consumption by a building's cooling system at times of high solar irradiation, which also correspond to high levels of PV power generation.

critical for adequate representation of the future. For example, the data may pertain to a time interval which is insufficiently short for capturing the stochasticity of the underlying processes.

The recognition of the complex nature of data, with all of its positive and negative aspects, is one of the critical and distinguishing components of the methodology of this work. Accordingly, we develop data-driven methods which partially relax the assumption that the historical data perfectly represents the future, and thus increase the effectiveness of the resulting studies.

Another distinguishing methodological consideration of this work is the recognition of the plurality of methods that may be used to achieve the same objective. Accordingly, we develop and compare multiple methods for studying each system problem. Each method may have distinct strengths and weaknesses. Consideration of multiple methods allows us to choose the best one or combine the advantages of several methods into a single joint approach.

Finally, the consideration of the *practicality* of the result is also one of the methodological premises of this research. Following this premise, we ensure that our assumptions are practically feasible. The practical limitations depend on the context and manifest themselves as restrictions on the computation time, the availability of data, and others. Abiding by these limitations makes the results of the research applicable to real systems.

# Chapter 2

# Terminology and Related Work

## 2.1 System Components

This dissertation focuses on PV-storage systems installed in a building or home, including off-grid systems that have no additional power sources, and grid-connected systems that have a bi-directional connection to a larger electricity grid. In this section, we discuss the basics of the main components of our system: solar PV panels, the load, the battery, and the grid, if present.

### 2.1.1 Solar PV Panels

PV panels are composed of many PV cells, which convert solar energy into electricity. The panels are typically either mounted to building rooftops or ground-mounted. The mounts may be either fixed or track the sun for increased generation. The generated power follows a diurnal pattern and is heavily dependent on the panel temperature and the solar radiation, which are dependent on location, season, and cloud cover. The power capacity of the panels is usually expressed as kilo-Watt peak (kWp, shortened to kW), which is the amount of power they generate when exposed to a solar irradiance of 1000 W/m$^2$ at a temperature of 25°C.

### 2.1.2 Battery

Charging (accumulation of energy obtained from external sources), storing (holding energy inside the battery), and discharging (providing energy to external consumers) are the three

main functions of a battery.

In this work, we focus on electrochemical batteries, specifically those which use Lithium-ion (Li-ion) technology. Li-ion batteries lie at the heart of many modern devices and systems, including smartphones, electric vehicles, grid-scale energy storage, and, increasingly, in buildings.

Li-ion batteries have seen a huge rise in popularity during the last decade. These batteries have many desirable properties, including high energy density, high efficiency, and fast charging/discharging capabilities. There are many variations in the chemical materials used in a Li-ion cell, and each variation results in a cell with different properties. For example, Lithium-Titanate cells have long life spans [103], while Lithium-Ferrous-Phosphate cells can be discharged at very high power [2].

The price of Li-ion batteries is expected to continue to fall over the next 20 years [124], which favours many applications where storage is a necessity or could be useful. Some key applications are electric vehicles, grid regulation, and the integration of energy storage with renewable energy sources.

A battery is typically composed of many Li-ion cells, which are arranged in a circuit and managed by a Battery Management System (BMS). The BMS limits the use of the battery within a safe operating range, with constraints on the terminal voltage, charging/discharging current, and operating temperature. A fraction of energy is lost upon conversion from electrical to electrochemical energy during charging, as well as upon conversion in the reverse direction during discharging, with a typical round-trip efficiency of 80-95%. The amount of energy that can be stored in the battery at any given time is expressed in kilo-Watt hours (kWh). For a building, a typical battery size may be in the range of 1-100 kWh. Further details on battery characteristics may be found in the discussion on battery models in Section 3.1.

### 2.1.3 Load

The electrical load on the system is the sum of the demands of the electrical appliances currently in use. For a residential building, major appliances include air conditioners, heaters, refrigerators, lighting, and electronics such as televisions and computers. In this work, we consider the load to be inelastic, in the sense that it follows a fixed stochastic trajectory which may not be modified. Some modern appliances such as washing machines and dishwashers allow for flexible scheduling of their operation to coincide with times of cheap grid pricing or low demand. These are not considered in this thesis.

### 2.1.4 Grid

The electrical grid is a large network for delivering electricity from producers to consumers. From the perspective of a building, in normal circumstances, we assume that electricity can always be purchased from or sold to the grid as needed. An electricity utility company sets the price of electricity for consumers in the form of a *pricing scheme*. There are thousands of variations for grid pricing schemes; the common parameters in those schemes include time of day, day of the week, season, peak power demand, and aggregate demand. For grid connections to commercial and residential buildings, we study the following two common pricing scheme structures:

- **Time-of-use pricing**: In this scheme, the price of electricity varies depending on the time of the day on a fixed schedule with two or more pricing periods. For example, in Ontario, there are off-peak, mid-peak, and peak pricing periods, each with their own per-kWh electricity price; the schedule of each period and their corresponding prices are updated on an annual or semi-annual basis, and typically correlate with the levelized cost of electricity production in the grid.
- **Peak-demand pricing**: The price of electricity is time-invariant, but with high penalties for exceeding a given peak power threshold.

In many countries, including Canada and the United States, there are programs in place for the grid to purchase electricity generated by local PV panels. There are three common price structures:

- **Net-metering**: Energy transferred to the grid is used to offset the energy consumed from the grid, thereby lowering the electricity bill of the system owner.
- **Feed-in tariff**: Energy is sold to the grid at a fixed rate, often guaranteed by a long-term contract. Feed-in tariff plans were originally used as a subsidy for PV panels, with rates that ensure a positive return-on-investment. The rates have declined alongside the declining cost of PV panels to below even the cost of grid electricity.
- **Value-of-solar tariff**: The price for selling PV energy to the grid is calculated to reflect the value of the benefits provided by distributed solar feed-in to the electricity grid operator [158].

## 2.2 Related work

Energy systems are a widely-studied area of research. At the start of each of the next three chapters, we survey the literature corresponding to each research problem. This section gives a general overview of how the problems in energy systems have been addressed (the methods), and the goals that the researchers have attempted to achieve (the objectives).

### 2.2.1 Common Methods

Existing work on energy system analysis may be grouped according to the method. In our work, three broad methods are considered: 1) simulation of the evolution of the physical system, denoted *Simulation*, 2) Mathematical programming to determine optimal sizing and operation, denoted *Optimization*, 3) solving a set of equations that mathematically describe the state and state evolution of the system, denoted *Analytical Methods*.

System simulation involves estimating the evolution of system state over time while keeping track of relevant system metrics. Often, the simulation is performed multiple times with modified parameters to determine the effect of parameter variations on the relevant metrics. For example, Wang et al. [169] developed a simulation for microgrids with renewable energy sources to study their performance under realistic operating conditions.

Mathematical optimization methods consist of finding an optimal value from a defined set or range of values that represent feasible system states. The optimization problem is commonly expressed as a *mathematical program*, with a function for calculating the optimization objective, and a set of constraints on the function. For example, Solanki et al. [150] formulate a mathematical program as part of an online optimization algorithm for the control of a microgrid with energy storage, conventional power generators, renewable energy sources, and deferrable loads. A comprehensive survey of optimization methods for renewable energy applications can be found in [18].

Analytical methods are those which define the system through a set of equations and mathematical operations, where the solution of the equations represents some relevant system state or parameter which is of interest to the researcher. Often, the solution is closed-form or can be evaluated efficiently using some basic algorithmic operations. For example, Wang et al. [172] extend stochastic network calculus methods to asses grid power supply reliability in the presence of renewable energy sources supported by batteries.

This thesis focuses on simulation and optimization methods. Complementary work on analytical methods can be found in reference [74, 75] where the formulation has been worked out by our collaborator, with an implementation by the author.

### 2.2.2 Common Performance Objectives

A commonly defined objective of energy system analysis studies, especially for those systems which include PV panel and battery components, is to minimize the monetary cost associated with the system. The high price of these components being the primary roadblock to their widespread adoption may explain this objective. The cost of these systems may be broken down into two components: one-time, up-front capital expenses (capex), and operational expenses (opex). Accordingly, the corresponding system engineering problems are:

- **System design (reducing capex)**: given an operation setting, the study of system design centered around minimizing the size (and thereby cost) of system components while meeting some system objectives or simply maximizing system return-on-investment.
- **System operation (reducing opex)**: given a system, the study of system operation focus on maximizing the system's value, given a particular system size, for example by reducing the cost of meeting the energy demands for the system owner.

Optimal system design and operation are coupled: the size of the system affects the optimal operation, and the choice of operating policy affects the optimal design. This coupling is practically impossible to break. Consequently, the problems of optimal design and operation are ideally studied together. However, when considering certain limitations, such as the fact that we are usually unable to predict the future operating environment with perfect accuracy, optimal design and control may not be a practical goal. Consequently, this coupling is often broken, at the cost of their joint optimality, and the problems are studied separately as:

- Given a practical operating policy, what system design will offer relatively low capex and opex?
- Given a system design, what practical operating policy will achieve relatively low opex?

System application, i.e., the purpose for which the system is being used, is another cost-related consideration. For example, the system may be used to directly reduce the owner's electricity bill, or to generate revenue by participating in a day-ahead electricity market as part of a virtual power plant [123], or to earn credits in a peer-to-peer energy

sharing cooperative [105]. *Stacking* multiple applications on one system has the potential to increase system value [3, 42]. The complexity associated with balancing the sometimes competing objectives of multiple applications has ramifications on the complexity of system design and operation problems. In our work, we treat the system application, as well as the choices of specific PV panel and battery technology, as being a fixed part of the sizing or operation problem.

# Chapter 3

# Energy System Modelling

Models of batteries, which estimate battery energy content after a series of charge and discharge operations, play a central role in the optimal design, analysis, and operation of battery-based systems. Many battery models have been developed in the past, with varying degrees of computational complexity, ease of use, and fidelity [143]. Some models aim to simulate the electrochemical processes within a battery, or emulate battery behaviour using electrical circuits; these approaches are highly accurate, but are computationally complex and therefore not suitable for use in optimization or large-scale simulation studies.

In this chapter, important characteristics of Li-ion batteries are discussed in Section 3.1. Existing models are compared with respect to a set of criteria for models that are well suited to system simulation and optimization studies in Section 3.2. This is followed by the development of new models for simulation in Section 3.3 and models for optimization in Section 3.4. Finally, we discuss limitations and propose potential directions for further developments of this theme in Section 3.5.

Note that the models developed in this chapter are not limited to PV-storage system analysis. They are application-agnostic and may be used in the study of any system featuring Li-ion batteries. Materials in this chapter have been published in [81, 78, 80]

## 3.1 Battery Characteristics

A typical battery system is composed of one or more cells and the associated battery management system (BMS), as shown in Figure 3.1. Power from a source flows to the BMS, which charges the cells at a particular charge current; the product of the cell voltage

Table 3.1: Battery Model Notation

| *Name* | *Description* |
|---|---|
| $\alpha_c$ | Maximum charging current |
| $\alpha_d$ | Maximum discharging current |
| $a_1$ | Lower energy limit function; depending on the model, this could be a constant approximation $\bar{a_1}$, a function of the power $a_1(P)$, or an interpolation of empirically derived points |
| $a_2$ | Upper energy limit function; depending on the model, this could be a constant approximation $\bar{a_2}$, a function of the power $a_2(P)$, or an interpolation of empirically derived points |
| $\eta_c$ | Charging efficiency; depending on the model, this could be a constant approximation $\bar{\eta_c}$, a function of power $\eta_c(P)$, or a function of the applied current and battery voltage $\eta_c(I, V)$ |
| $\eta_d$ | Discharging efficiency; depending on the model, this could be a constant approximation $\bar{\eta_d}$, a function of power $\eta_d(P)$, or a function of the applied current and battery voltage $\eta_d(I, V)$ |
| $\Delta_E$ | Difference in energy content |
| $E$ | Battery energy content |
| $I$ | Current |
| $M$ | Voltage function, relates battery energy content and applied current to battery voltage |
| $n_c$ | Number of cells in battery |
| $P$ | Power |
| $Q$ | Charge (Coulomb) content |
| $R_{ic}$ | Internal resistance during charging |
| $R_{id}$ | Internal resistance during discharging |
| $u_1$, $v_1$ | Parameters to linear approximation of lower energy limit $a_1$ |
| $u_2$, $v_2$ | Parameters to linear approximation of lower energy limit $a_2$ |
| $V$ | Voltage |
| $V_{nom,c}$ | Nominal voltage during charging |
| $V_{nom,d}$ | Nominal voltage during discharging |
| $x_{00}$, $x_{01}$, $x_{10}$ | Parameters to planar approximation of $M$ function |

and the charge current is the injected power. Symmetrically, the product of the cell voltage and the discharge current is the power drained from the cells by the BMS for discharge. The cells are arranged in a circuit with a specific topology, composed of parallel and series

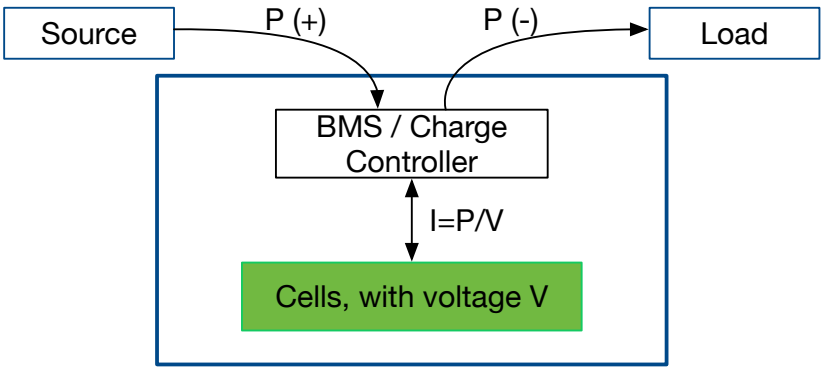


Figure 3.1: Battery system components

components; the structure affects the overall voltage of the battery as well as the maximum current it can handle, but does not affect the maximum power rates and energy capacity.

A Lithium-ion cell is an electrochemical system with three basic elements: the positive electrode (cathode), the negative electrode (anode), and the electrolyte, as illustrated in Figure 3.2. The cathode is a Lithium compound, the anode is typically carbon-based, and the electrolyte works as a conductor that allows Lithium ions to move between the anode and cathode with little resistance. There is also a separator material between the anode and cathode that prevents their direct contact but allows $Li^+$ ions to pass through. When a positive current is applied to the battery, $Li^+$ ions flow from the cathode to the anode and charge the battery; when a negative current is applied, the $Li^+$ ions flow from the anode to the cathode. This reaction causes some side-reactions that have cumulative effects over time and use, and limit the lifetime of the battery.

The electric potential difference between the positive and negative battery electrodes is the *voltage* of the battery. The BMS confines the voltage of each cell to a range of acceptable values to avoid damaging the cells. We will refer to the endpoints of the voltage range as $V_{min}$ and $V_{max}$. Specifically, the BMS protects the cells from being damaged through overcharging or undercharging by restricting charging when the voltage reaches $V_{max}$, and restricting discharging when the voltage drops to $V_{min}$. The BMS also prevents cell damage caused by exceedingly high charging or discharging power flows via reduction of the power to the levels determined by the manufacturer. Finally, the BMS may restrict the charge level of the cells within a "depth-of-discharge" limit, to increase the service life of the battery [131].

A battery is often described using a few key parameters: the number of cells, the nom-

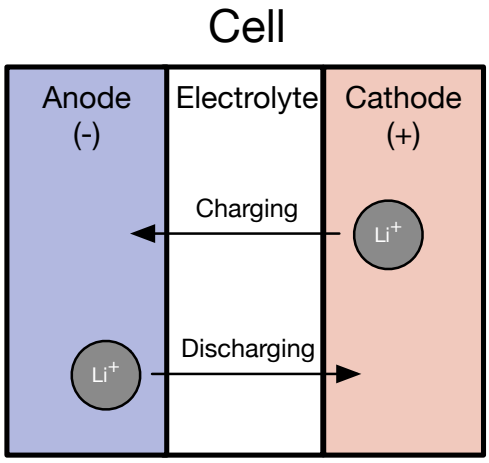


Figure 3.2: Illustration of the internal structure of a Lithium-ion cell.

inal voltage[1], the maximum charging and discharging rates, and the number of charge-discharge cycles the battery may be subjected to before it is considered to be worn out. Li-ion batteries can be manufactured using a variety of different chemistries, and each chemistry has distinct properties. For example, Lithium-Titanate (LTO) cells have a nominal voltage of 2.3 V, last for over 10,000 cycles, and may be charged and discharged relatively quickly. In contrast, Lithium-Ferrous-Phosphate ($LiFePO_4$) cells have a nominal voltage of 3.2 V, a cycle life of approximately 5000 cycles, and also support high charging and discharging rates. Similarly, Lithium-Nickel-Manganese-Cobalt-Oxide (NMC) cells have a nominal voltage of 3.7 V, a cycle life of approximately 2000, and support relatively low charging and discharging rates. Many of the battery's properties are affected by temperature and level of battery degradation; we do not consider these factors in our models and defer the discussion of these limitations to Section 3.5.

Batteries have many non-linear dynamic properties, including the voltage trajectory, charging/discharging efficiencies, and energy content limits. To help explain these, we will refer to the battery specifications (spec) sheet of a Leclanché LTO cell [103]. The spec provides 'reversible capacity' curves, such as the ones in Figure 3.3, which show the trajectory of the voltage during charging or discharging at a constant applied current. In a typical spec sheet, charging and discharging currents are expressed in terms of a *C-rate*, where 1 C is defined as the current needed to fully charge or discharge the nominal

[1]Nominal voltage refers to the average voltage that is typically the midpoint of the voltage range of the battery.

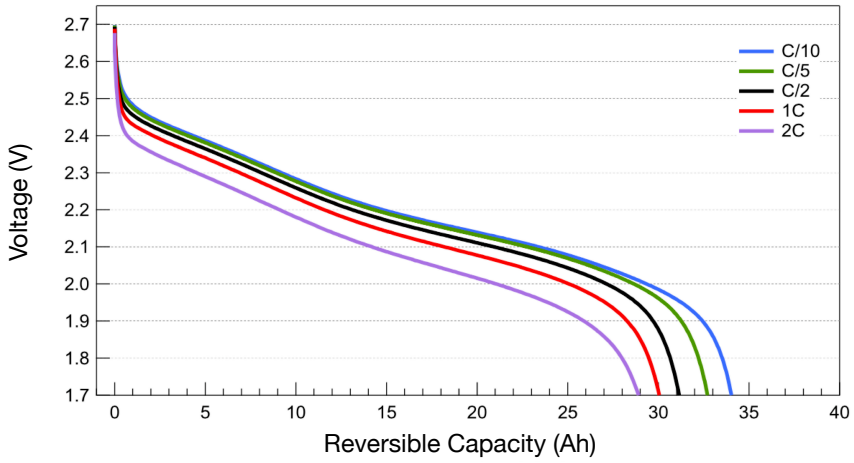


Figure 3.3: Voltage vs. charge content for different discharge rates [103].

capacity of the battery (which is specified in the spec sheet) in 1 hour. For the LTO cell, 1 C corresponds to 30 A, because the nominal charge capacity is measured to be 30 Ah.

The trajectory of the terminal voltage as the battery is charged and discharged is non-linear, as shown in Figure 3.3 for discharging. The voltage trajectory depends on the applied current. The relationship between battery voltage $V$, applied current $I$, and corresponding applied power $P$ is:

$$P = IV. \tag{3.1}$$

A given $P$ can correspond to many different combinations of $I$ and $V$ values. The true value of $V$ depends on the applied current, as well as on the energy content.

The charging energy efficiency of the battery is the fraction of energy that is stored in the battery after losses in the energy conversion process from electrical to electrochemical potential have occurred. Similarly, the discharging energy efficiency is the fraction of energy that is lost in the reverse conversion. Both charging and discharging efficiency are proportional to $I^2R$, where $R$ is the internal resistance of the battery. The internal resistance may be different for charging and discharging processes.

A 'nominal energy capacity', measured as a specific discharging rate (often 1 C) is often used to describe the energy capacity of a battery, i.e., the amount of energy that can be discharged from a full battery before the voltage reaches $V_{min}$. Battery spec sheets provide battery voltage curves as a function of 'reversible capacity,' expressed in units of Ah or Coulombs, which is the amount of charge added or removed from the battery. Since the

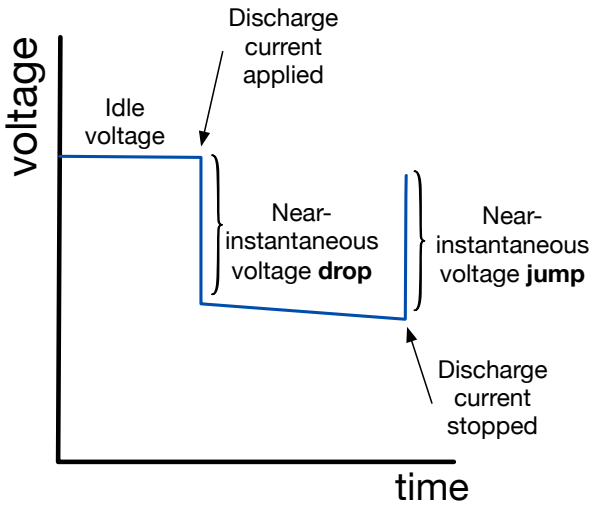


Figure 3.4: Illustration of voltage drop when an increased discharging current is applied to the battery.

curve is expressed in volts, and 1 Joule = 1 Volt-Coulomb = 1 Watt-second, the area under each curve in Fig. 3.3 indicates how much energy can be discharged from the battery at the given C-rate. It is clear that as the C-rate increases, the energy capacity of the battery decreases. Similarly, the area under each curve in Fig. 3.8 represents the energy needed to charge an empty battery at a given C-rate. We describe this behaviour in greater detail next.

### Charging and Discharging Behaviours

Accurate modelling of the dynamic behaviours of a Li-ion battery requires an understanding of the effects of different rates of charging and discharging rate on the voltage. When the battery is idle, i.e., the applied power is 0, its voltage is a good indicator of the amount of energy it contains. An idle voltage equal to $V_{max}$ indicates that the battery is fully charged. However, when a discharging current is applied, the battery voltage instantaneously drops, only to jump back up when charging is stopped, as illustrated in Figure 3.4. The magnitude of the voltage drop is approximately proportional to the current. Similarly, there is a voltage jump during charging that is roughly proportional to the magnitude of the charging current.

To prevent the battery voltage from exceeding $V_{max}$ during the charging process, espe-

cially when the battery's energy content is high, the BMS must limit the charging current. This is the basis for widely-used protocol for quickly charging a Li-ion battery: Constant Current - Constant Voltage (CC-CV) charging [185], which has two phases. In the CC phase, the battery is charged at a high current until it is nearly full, which causes the voltage to reach the upper limit. Subsequently, in the CV phase, the current is gradually reduced while maintaining the voltage at the upper limit.

The dynamics of the CC-CV charging protocol could be illustrated through an analogy with filling up an empty glass with water from the tap. Initially, while the glass is mostly empty, a strong water stream quickly fills up the glass without causing an overflow. However, the strong stream creates turbulence of water in the glass, which could cause water to spill out even though the glass is not yet full. This is why, when the glass is almost full, the incoming stream must be gradually reduced to a trickle, thereby reducing turbulence and completing the filling of the glass without spills.

## 3.2 A Survey of Battery Models

There is a rich and growing corpus of work on the topic of modelling energy storage. These models take as input the power or current applied to the battery as well as the prior battery state, and estimate the resulting voltage (V), energy content (Wh), and charge (Ah), which describe the new state of the battery. The most sophisticated examples are based on modelling the internal chemical state of the battery, or modelling the battery as an equivalent electrical circuit, while others attempt to describe only high-level dynamics via simple mathematical expressions. There are hundreds of models in the literature; we review several representative examples and focus on their common features.

Detailed *electrochemical* models attempt to simulate the chemical processes that occur in the cells, such as the movement of $Li^+$ ions through the cell. These models typically focus on the cells and do not describe the actions of the BMS. The model parameters include the physical properties of the cell such as the thickness of the separator and electrodes, and chemical properties such as the electrolyte concentration and conductivity [90].

*Equivalent-circuit* models aim to emulate the behaviour of a battery by using electrical circuits [60, 63]. These models are not meant to express the physical properties of the cell, and hence many of their parameters do not have a physical meaning and must be calibrated via curve-fitting or experiments.

*Simplified mathematical* models belong to the class of models that aim to describe electrochemical battery behaviour at a high level. Many of these models are based on

Shepherd's model for electrochemical batteries [146], which calculates battery voltage as a function of terminal voltage, open-circuit voltage, and internal resistance characteristics along with state of charge and current. Shepherd's approach has been used as a base for more sophisticated models [134, 175], and also forms the basis of the generic battery model for estimating battery voltage and charge level in Matlab's Simulink package [166, 165].

In developing our models, we take the perspective of a systems analyst who desires a battery model which is both accurate and easy to use. A model that is well suited to system studies is expected to satisfy the following conditions or criteria:

1. *Accuracy.* The modelled battery's behaviour should conform to the behaviour of a real battery. Specifically, the error accumulated in the modelled energy content after a series of charge/discharge operations should be minimized.

2. *Acceptable computational time.* The model should compute an output within an acceptable time, given the computational power available to the researcher.

3. *Optimization-friendly.* Where possible, the model should be composed of analytic expressions. This is a necessary condition for a model that is suitable for optimization via mathematical programming.

4. *Reasonable calibration effort.* Model calibration should involve a minimum amount of effort, ideally requiring only data that is readily available, such as the data found in a battery datasheet.

5. *Power input.* The model should use power, rather than current, as input. Working with power is preferable because power is conserved as it flows through a power system, as opposed to current which may go through many transformations. Using a model that takes current as input implies additional and thus undesirable effort to explicitly model the transformations.

6. *BMS-aware.* The model should include the two main components of the battery, i.e., the cells and the BMS. The constraints typically imposed by a BMS should be integrated into the model. A model satisfying this criterion of completeness is preferable to an otherwise incomplete model of cells relying on some external BMS model.

This set of conditions is specifically chosen to highlight the distinctions of existing models from the point of view of their suitability for simulation or optimization studies, and we use them as a guide for developing our models. In this section, we review the existing literature on battery models with respect to these criteria.

### 3.2.1 Model Complexity

In the context of this work, the complexity of a model refers to the computational resources needed to compute a model output.

All other things being equal, less-complex models are preferred. Ideally, not only should a model be simple, but it should also minimize the complexity added to an algorithm, such an as optimization solver, that uses it. A model that calculates the change in the battery's energy content explicitly, rather than as a fixed point of a set of coupled equations is clearly simpler. Similarly, a model based on linear equations and constraints is convex and hence simpler to solve as part of a mathematical program, compared to a model with non-convex elements.

Electrochemical models are among the most accurate, and unfortunately also the most complex models. They express dynamic battery behaviour in terms of the chemical properties of the battery materials, in particular, the anode, cathode, and separator components. Electrochemical models are generally composed of differential and partial-differential equations describing the rates of Lithium-ion diffusion and their concentration [148].

Another type of model is the equivalent-circuit model, which emulates battery dynamics using an electrical circuit. Even though these can be linear circuits, it has been shown that at least one resistance-capacitance (RC) pair is needed in the circuit to predict the terminal voltage of the battery with acceptable accuracy, or with errors of $< 5\%$, across a range of test conditions [63, 120, 57]. RC pairs introduce differential equations into the model, making it complex to solve.

These electrochemical and equivalent-circuit models (with RC components) can be used in simulation studies, but they cannot be integrated into a mathematical programming framework due to their use of non-analytic equations, such as differential equations. On one hand, the idea of approximating the non-analytic aspects of these models with analytic functions is feasible [101]. On the other hand, such approximation may or may not result in a suitable model from the perspective of a mathematical programming framework.

The accuracy of approximations and the amount of computational resources required for the model are important aspects. Hence, great care must be taken in approximating the non-analytic, non-convex, and non-linear components of these models to make them suitable for integration into mathematical programs.

### 3.2.2 Parameter Calibration

Parameter calibration is the process of determining the value of model parameters. Calibrating the parameters for many of the developed battery models is a challenge. For example, equivalent-circuit models require pulse-current charge/discharge voltage measurements from which model parameters are extracted ([29, 61, 64, 65, 67, 89, 91, 182]), while neural network approaches [25, 37] require data sets of battery measurements for training the network. These data sets are difficult to obtain, which makes the models challenging to use in practice. Obtaining these parameters of the model requires experiments with the battery on specialized equipment, which increases the complexity of any study that uses the model or limits it to modelling batteries for which this data is already available.

Given this complexity, we prefer models that can be calibrated using only readily available data, in particular, the data typically found in the spec sheet released by a manufacturer of the battery cells[2]. The approach of basing the calibration of model parameters on spec sheets is reasonable; this information is authentic, verifiable, and in principle should capture the main characteristics of the battery. This approach is an alternative to running experiments for the sake of extracting the parameters from battery measurements. In essence, a spec sheet is a summary of the experiments conducted by the cell manufacturer.

The contents of a spec sheet vary by manufacturer; however, a spec sheet typically includes information about nominal capacity, internal resistance, and voltage vs. ampere-hour reversible capacity curves for battery charging and discharging at different currents. Thiruganam *et al.* [162] propose a model that can be calibrated using only the information found in the spec sheets released by battery manufacturers such as EIG, Sony, Panasonic, and Sanyo. Their parameter calibration approach is complex, involving the use of a genetic algorithm to fit the polynomial function parameters of their model to the voltage curves from a spec sheet. A simpler and more widely used *spec-based* model, where model parameters are calibrated using only the spec, for simulation was proposed by Tremblay *et al.* [166] for four different battery chemistries: Lithium-ion, Lead-acid Nickel-Cadmium, and Nickel-Metal-hydride. A later publication [165] describes improvements to the model equations that are specific to each battery chemistry. Tremblay's model is the default battery model in the widely used Matlab Simulink software package.

An important aspect of all battery models is that the model equations offer approximations to real physical processes. The parameters of these equations may be calibrated to optimize model accuracy to a specific battery operating range; the narrower the range, the closer the approximation can be calibrated to resemble the real process. The system

[2]Examples of spec sheets can be found at http://category.alldatasheet.com

application is, therefore, a crucial parameter of the model calibration process; modelling an application that uses only a small fraction of the battery's full operating range should be done with a calibration that focuses on the accuracy only within that range. We discuss this in greater detail in Section 3.4.2.

### 3.2.3 Model Input

The input to all electrochemical and equivalent circuit models that we have seen in the literature is the applied current, which is crucial for accurately estimating the dynamic behaviour of a battery. However, when conducting a power flow simulation as part of a system optimization study, calculating the current applied to the battery can be challenging because of the transformations in the current that may occur between the power source and the battery, or between the battery and the load. Furthermore, we don't know the applied current until we calculate the voltage of the battery, due to the coupling between battery voltage, power, and current described in Eq. (3.1). For these reasons, a battery model that takes power as input is preferred.

### 3.2.4 Battery Management System

Based on our analysis of the literature, we find that the majority of existing battery models focus on modelling the cells, without taking into account the constraints that would be imposed by a BMS. One possible explanation is that the cells exhibit the majority of the complex and interesting behaviours of the battery, and so has been the source of research problems. In contrast, BMS operation is usually straight-forward. However, the BMS is an integral part of the battery. It prolongs battery life and ensures safe operation. For this reason, this crucial component should be included in battery system models.

### 3.2.5 Basic Model

The computational effort required to solve a mathematical program is proportional to the number and complexity of equations and model constraints in the program [132]. The simplest programs to solve are linear convex programs (LP), which have reasonable solving times even when the program has hundreds of thousands of constraints [47]. In studies requiring an optimization framework, a widely-used battery model is linear (for examples, see [85, 59, 111, 12, 30]). While the specific model in each work may differ in some details, the structure of the model remains essentially identical, and we refer to it as the *basic*

model. This model uses power as input and has linear and constant approximations to many of the non-linear dynamic characteristics of a Lithium-ion battery.

In our formulation of the basic model, we use the following mathematical notations. Given the energy content, $E(k-1)$, of a battery in time-slot $k-1$, and an applied power $P(k)$ which is constant over time slot $k$, the model calculates an estimate of the energy content $E(k)$ at the end of the $k^{th}$ time slot. Three basic properties are modelled: 1) the charging and discharging inefficiencies ($\eta_d$ and $\eta_c$, respectively), 2) the upper and lower energy content limits enforced by the BMS ($a_1$ and $a_2$, respectively), and 3) the charging and discharging current limits also enforced by the BMS($\alpha_c$ and $\alpha_d$, respectively), which are converted into power limits by multiplying with the corresponding nominal cell voltage ($V_{nom,c}$ and $V_{nom,d}$). The values of these parameters can be obtained from the spec sheet of the cell. For a battery system comprised of $n_c$ cells, the battery energy and power limits scale linearly with $n_c$, i.e., a battery composed of two cells has twice the energy capacity and power limits as that of a single-cell battery.

The typical formulation of the basic model, for example in [48], for a battery with $n_c$ cells is given below. The piece-wise Eq. (3.3) is often expressed in linear form by splitting $P(k)$ into separate charging and discharging variables, say $c(k)$ and $d(k)$ respectively. It is physically impossible for a battery cell to be charged and discharged simultaneously, which yields the following non-linear constraint $c(k) \cdot d(k) = 0$. It has been shown in [51] that this constraint can often be relaxed, since simultaneous charging and discharging is rarely optimal due to efficiency losses, which makes the model linear.

**The basic model ($n_c$ cells):**

$$E(k) = E(k-1) + \Delta_E(k) \tag{3.2}$$

$$\Delta_E(k) = \begin{cases} \eta_c P(k) T_u & : P(k) \geq 0 \\ \eta_d P(k) T_u & : P(k) < 0 \end{cases} \tag{3.3}$$

$$n_c \alpha_d V_{nom,d} \leq P(k) \leq n_c \alpha_c V_{nom,c} \tag{3.4}$$

$$n_c a_1 \leq E(k) \leq n_c a_2 \tag{3.5}$$

Note that $\eta_c$, $\eta_d$, $a_1$, and $a_2$ are all constant, and the accuracy of the model is heavily dependent on proper calibration of these parameters. However, we have noticed that it is common practice to calibrate the model without considering the following: 1) the operating range of the application (discussed in Section 3.4.2), and 2) the characteristics of the specific battery chemistry being used, which can be found in a spec sheet. In the literature on energy system optimization involving a battery, typical values used are a discharging and

charging efficiency in the range of 90-95%, $a_2 = B$ where $B$ is the nominal energy capacity, and $a_1 = 0$ or a depth-of-discharge limit of $0.2 \cdot B$ or similar value to extend battery service life.

The basic model has its advantages and disadvantages. It has an advantage over many other models in its clarity and simplicity of calculations. On the other hand, the potential loss of accuracy as a result of the approximations to the complex dynamic behaviours of the battery is a disadvantage. We are unaware of any published works (other than our own) that thoroughly evaluate the basic model's accuracy, explain how to calibrate it properly, and explore how it affects the results of the optimization studies that use it. We believe that these issues are resolved in the sections below.

A significant shortcoming of the basic model is that, in a real battery, both the efficiencies and energy content limits of a battery are dependent on the current being applied and on the battery voltage. The basic model does not have an internal representation of the voltage and current, which hinders its accuracy. Our attempt to resolve this problem is described in Sections 3.3 and 3.4, where we develop new models and where we compare the accuracy of the newly developed models with the accuracy of the basic model.

## 3.3 Model for Simulation

In this section, we develop a model that we call the power-based integrated model (PI model). Our modelling objective is to capture the complex characteristics of Li-ion batteries, while at the same time relying only on the data typically available in a spec. In the process of achieving this objective, we used the conditions described in Section 3.2 as a guide.

We find that to accurately model the non-linearities inherent in a battery, it is necessary to implicitly model the battery terminal voltage resulting from a power injection. In doing so, we introduce non-analytic components to the model. These non-analytic components cause the model unsuitable for optimization studies using mathematical programming, though the model remains suitable for simulation studies. In Section 3.4, we develop new battery models for optimization using the PI model as a starting point.

The PI model has been implemented as a Matlab system block compatible with the Simulink package. The details of the implementation are provided in Section 3.3.4.

### 3.3.1 Formulation

The primary goal of the PI model is to accurately estimate the voltage of the battery as it is subject to charging and discharging. The PI model attempts to capture the relationship between battery terminal voltage $V$, applied current $I$, and energy content $E$. This relationship exists in a steady state when the temperature and degradation level of the cells is held constant. A typical battery spec contains voltage curves measured across full charge-discharge cycles at different constant currents; these curves contain enough information to reasonably estimate the steady-state relationship between voltage, current, and energy content.

The PI model incorporates the following high-level characteristics of Lithium-ion batteries:

- **Charging and discharging penalties** We model battery charging and discharging inefficiencies as penalty functions of the current and voltage $\eta_c(I, V)$ and $\eta_d(I, V)$. $\eta_c(.)$ represents the fraction of applied power that is stored in the battery when it is being charged at power $P$, i.e., $(1 - \eta_c(.))P$ is the power loss penalty due to charging inefficiency. Symmetrically, $\eta_d(.)$ represents the power that needs to be discharged to obtain power $P < 0$ from it, i.e., the battery loses energy at a rate of $\eta_d(.)P$. Note that $\eta_c(.) \leq 1$ and $\eta_d(.) \geq 1$.

- **Charging/discharging rate limits** In our model, $\alpha_c$ denotes the maximum charging current, and $\alpha_d$ the maximum discharging current.

- **Energy content limits** A BMS limits the battery voltage to the range $[V_{min}, V_{max}]$. The amount of energy that can be obtained from a battery that is initially fully charged before the voltage reaches $V_{min}$ depends on the discharging current being applied [127]. Similarly, the charging current affects the amount of energy that can be stored in the battery before the voltage reaches $V_{max}$. In our model, instead of a permissible *voltage* range, we specify a permissible *energy content* range, using content limits that are functions of the current (which in turn depends on the charge/discharge power $P(k)$). Specifically, the upper limit $a_1(I)$ is the energy content at the point when the voltage reaches $V_{min}$ while being discharged with a current of $I$, and the lower limit $a_2(I)$ is the energy content at the point when the voltage reaches $V_{max}$ while being charged at a current $I$.

The following set of equations and constraints describes the PI model:

**The PI model:**

$$E(k) = E(k-1) + \Delta_E(k) \tag{3.6}$$

$$\Delta_E(k) = \begin{cases} \eta_c(I(k), V(k))P(k)T_u & : P(k) \geq 0 \\ \eta_d(I(k), V(k))P(k)T_u & : P(k) < 0 \end{cases} \tag{3.7}$$

$$V(k) = M(E(k), I(k)) \tag{3.8}$$

$$I(k) = \frac{P(k)}{V(k)} \tag{3.9}$$

$$\alpha_d \leq I(k) \leq \alpha_c \tag{3.10}$$

$$a_1(I(k)) \leq E(k) \leq a_2(I(k)) \tag{3.11}$$

Eqs. (3.6) and (3.7) indicate that the energy content at the end of time slot $k$ is simply the energy content at the end of time slot $k-1$ plus (resp. minus) the energy put into (resp. drawn from) the battery in time slot $k$ due to charging (resp. discharging). Eq. (3.8) maps the energy content and the charge/discharge current to the battery voltage and is discussed in greater detail in Section 3.3.1. The two internal variables $I(k)$ and $V(k)$ are *also* related through Eq. (3.9) which states that the power is the product of current and voltage. Eqs. (3.10) and (3.11) have been discussed above.

The PI model is therefore characterized by three variables (V, I, E), two parameters ($\alpha_c$, $\alpha_d$), and the five functions $\eta_c(I,V)$, $\eta_d(I,V)$, $a_1(I)$, $a_2(I)$ and $M(E,I)$. Given the energy content of the previous time slot $E(k-1)$ and a feasible input power $P(k)$, the model calculates the voltage $V(k)$, current $I(k)$, and battery energy content $E(k)$. In the following subsection, we will discuss Eq. (3.8) and how to use it. Section 3.3.2 explains how to calibrate these parameters and functions using the battery's spec sheet.

### The function $M$

Most battery spec sheets provide a family of curves that represent the battery's terminal voltage as a function of its 'reversible capacity'[3] when discharging or charging the battery at different rates. Fig. 3.3 and Fig. 3.8 (reproduced from [103]) show these functions for the Leclanché LTO cell. Although the spec sheet only provides this data for some C-rates, we can view these curves as representing a continuous surface. A key aspect of the PI model is to use a *mesh* function $M$ (a discretization of the underlying continuous surface) to map the battery energy content and charge/discharge current to the battery's voltage.

[3]This term is explained in Section 3.1

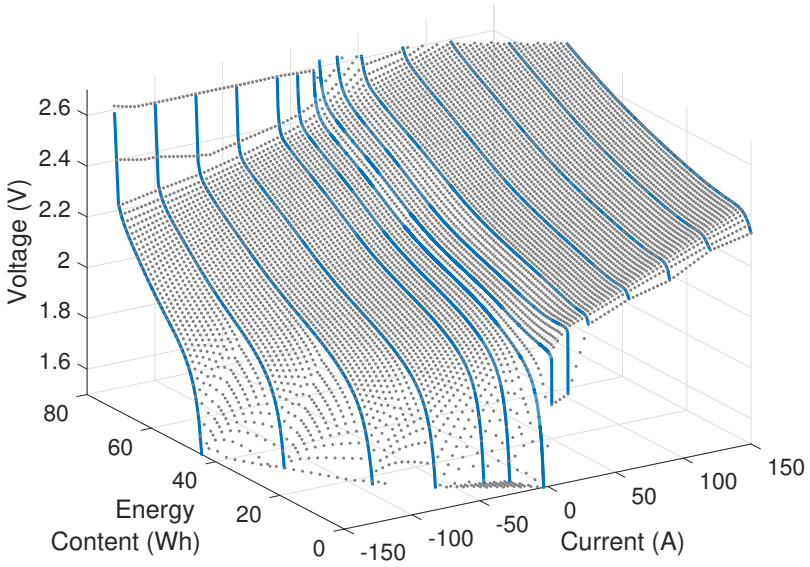


Figure 3.5: $M$ function, showing the terminal voltage as a function of energy content and charge/discharge current. The blue lines are the reversible capacity curves measured for an LTO cell [103]. The function, in grey, is a discretization of the underlying continuous surface.

This function, which can be derived and extrapolated from the spec sheet's reversible capacity voltage curves (see Fig. 3.3 and Fig. 3.8), captures the inherent non-linearities in battery voltage behaviour. An example can be seen in Fig. 3.5. Intuitively, $M$ represents all feasible combinations of energy content, charge/discharge current, and battery voltage.

We now discuss how to use the $M$ function to compute the change in battery voltage due to a certain amount of power injection or discharge $P(k)$. Note that the battery voltage has to be feasible (lie on the surface discretized by $M$) and consistent with the amount of power injected/discharged, i.e., $V(k) = P(k)/I(k)$. If $M$ had been defined in the space $< I(k), P(k), V(k) >$, this would correspond to finding the intersection, if it existed, between the curve $V(k)I(k) = P(k)$ and the surface described by $M$. This would be equivalent to finding the root of the function $M(P(k), I(k)) - P(k)/I(k) = 0$, which, since $M$ is an empirically-derived mesh function, can be solved numerically. However, $M$ is defined in the space $< I(k), E(k), V(k) >$ not $< I(k), P(k), V(k) >$. Thus, we need to define an auxiliary space, as discussed next.

Consider the subset of the $M$ surface obtained by setting $E(k)$ to the current energy content, b, which is described by:

$$\begin{aligned} S := \{< I, b, V >:\ & \alpha_d \leq I \leq \alpha_c, \\ & a_1(I) \leq E = E(k)^I \leq a_2(I), \\ & V_{min} \leq V = M(E, I) \leq V_{max}\}, \end{aligned}$$

where $E(k)^I$ is calculated using $I$ as the estimate for current using Eqs. (3.6) and (3.7). By construction, every point in $S$ satisfies all of the constraints of the model except for Eq. (3.9), i.e., $P(k) = V(k) \cdot I(k)$.

We then construct, in an auxiliary 2D space $< I, P >$, the line $y = P(k)$ (this is the given power injection/discharge) and the curve $y = VI$ for all points $(I, E, V) \in S$. The latter curve is obtained by enumerating all points in $S$, then computing, for each point, the corresponding pair of $(I, V \cdot I)$. The next step is to compute the intersection(s), if any, between these two curves. We do so by examining the points on the curve $y = V \cdot I$ in increasing order of $I$, searching for adjacent pairs of points $(I_1, V_1)$ and $(I_2, V_2)$ such that the sign of $P - V_1 \cdot I_1$ is different from that of $P - V_2 \cdot I_2$. The difference in sign indicates an intersection in between the two points, in which case we linearly interpolate between the points to get an estimate for the value of $I$ at the point of intersection. $E$ and $V$ values are then computed using this estimate.

Now, there are three possible cases (an example for each case is shown in Fig. 3.6):

**No intersection** There are no *feasible* voltage and current estimates for the given $P(k)$ and $E(k-1)$ (see Fig. 3.6(a)). In that case, the power is iteratively reduced by a small value and $S$ is recomputed for each value of $P(k)$ until an intersection is found. This value of $P(k)$ is the maximum feasible power that the BMS can allow.

**One intersection** Let that intersection point be $(I^*, P(k))$ (see Fig. 3.6(b)). In that case, $I^*$ is the estimate for $I(k)$ and $V(k) = \frac{P(k)}{I^*}$ the voltage estimate.

**Multiple intersections** The non-linearity of the surface makes it possible for multiple intersections to occur, as shown in Fig. 3.6(c). This case happens rarely for the Lithium-ion cell chemistries we have tested, and only occurred near the steep part of the voltage curve[4].

[4]The steep part of the voltage curve is seldom visited by many applications because of the practice of restricting the state of charge to preserve battery lifetime [113].

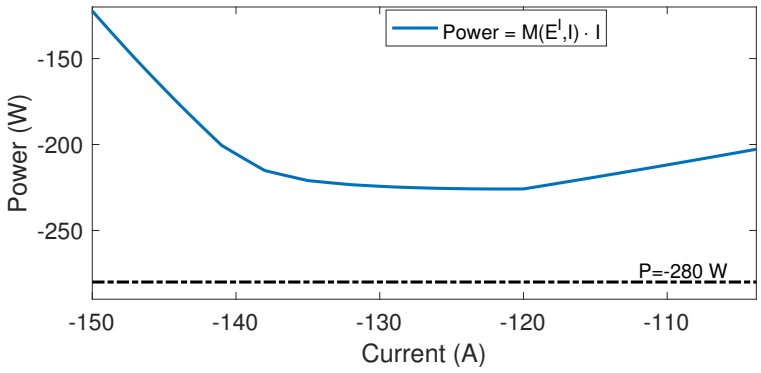


(a)

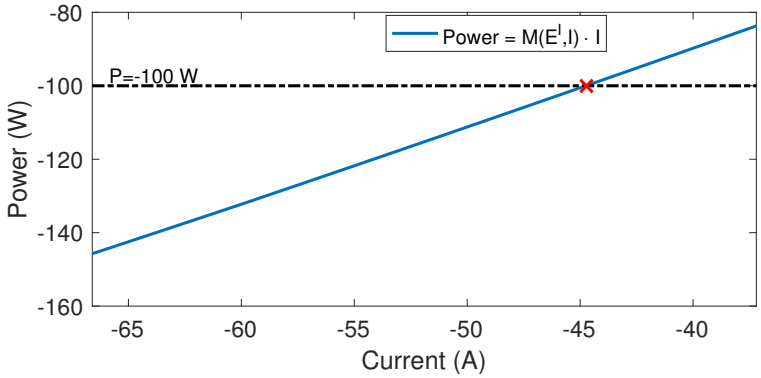


(b)

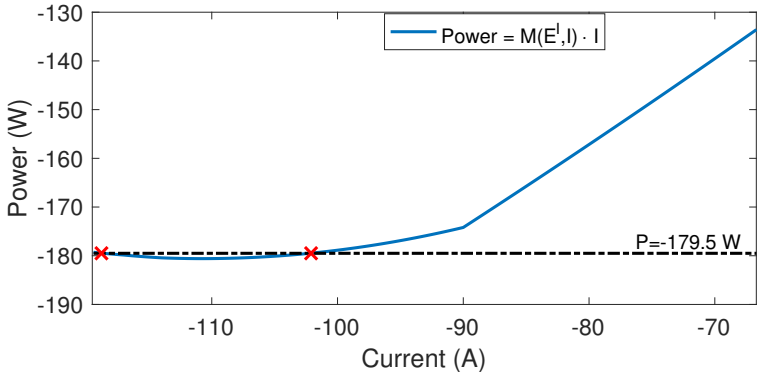


(c)

Figure 3.6: Three cases: No intersection (a), one intersection (b), and multiple intersections (c). A red cross indicates an intersection.

Table 3.2: Battery Specifications

| **Value** | **LTO** | **$LiFePO_4$** |
|---|---|---|
| Nominal Capacity (Ah) | 30 | 1.1 |
| $V_{min}$ (V) | 1.7 | 2 |
| $V_{max}$ (V) | 2.7 | 3.6 |
| Internal Impedance ($\omega$) | 0.002 | 0.05 |
| Max. charge (discharge) C-rate | 4 (4) | 4 (10) |

There are many potential approaches to choosing one of the multiple intersection points as our estimate. To preserve the continuity of voltage estimates over time, we recommend choosing the point yielding a voltage estimate which is the closest to the voltage value estimated for the preceding time step ($V(k-1)$).

Next, we explain how to calibrate our model.

### 3.3.2 Calibration

Calibrating the parameters of the model requires a battery spec sheet. We use the Leclanché LTO cell specifications document [103] as an example. Note that spec sheets may have varying amounts of information. The minimum information needed for calibrating the PI model is two voltage curves measured at different C-rates for both charging and discharging[5], the nominal capacity, the internal resistance, the voltage limits, and the charging and discharging rate limits.

Note that although the description of our model calibration process may appear complex, its inputs are easily obtained and the calculations are relatively simple. For example, although we present some equations as integrals, in practice we compute these numerically as a Riemann sum. Moreover, this calibration process has been automated as part of our Simulink implementation[6].

Table 3.2 shows the spec data (excluding voltage curves) for the LTO and $LiFePO_4$ cells.

[5]We require at least two curves to interpolate and extrapolate the M function for both charging and discharging processes.

[6]Matlab File Exchange link:
https://www.mathworks.com/matlabcentral/fileexchange/63078-lithium-ion-pi-model

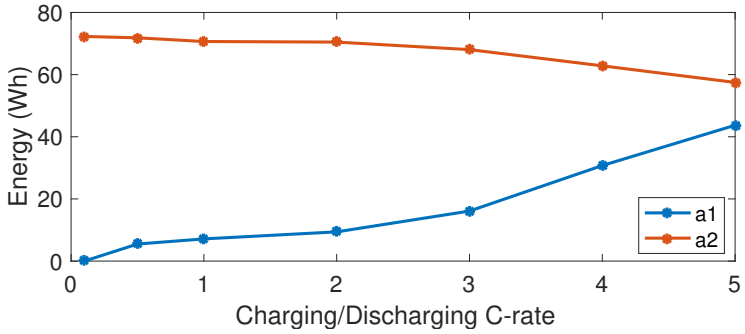


(a) $a_1(.)$, $a_2(.)$ energy content limits

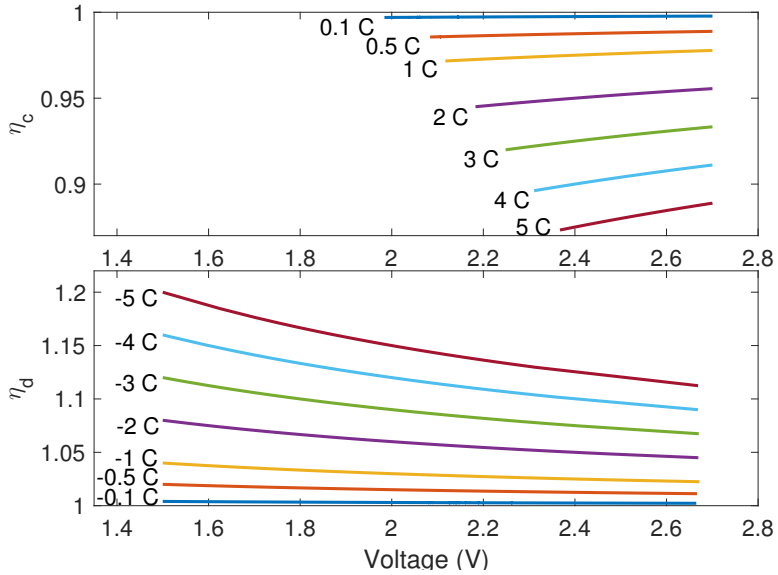


(b) Efficiency penalty functions $\eta_c(I(.), V(.))$ and $\eta_d(I(.), V(.))$

Figure 3.7: Parameters obtained for an LTO cell [103]. 3.7a): $a_1(.)$ and $a_2(.)$, points obtained from spec and linearly interpolated. 3.7b): $\eta_c(.)$ and $\eta_d(.)$, each line represents the efficiency at a different C-rate, as per Eqs. (3.12) and (3.13).

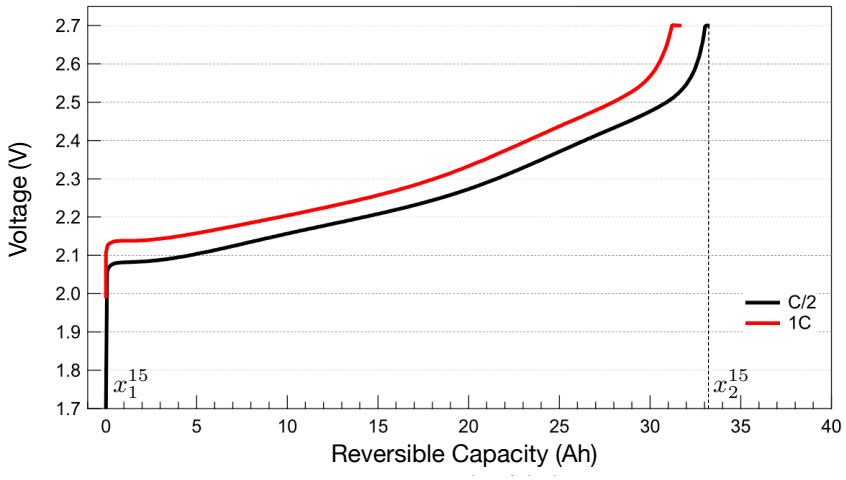


Figure 3.8: Voltage vs. charge content of LTO cell [103]. $x_1$ and $x_2$ for a C/2 charging current (15 Amperes) are labeled.

**Charge/discharge penalty functions:** $\eta_c(I,V), \eta_d(I,V)$

Let $R_{ic}$ and $R_{id}$ denote the internal impedance values (provided by the spec) during charging and discharging, respectively. Note that some specs provide only a single value $R_i$ for internal impedance; in that case, we assume $R_{ic} = R_{id} = R_i$. Then, it is easy to show that:

$$\eta_c(I,V) = 1 - \frac{IR_{ic}}{V} \text{ for } I > 0 \tag{3.12}$$

$$\eta_d(I,V) = 1 - \frac{IR_{id}}{V} \text{ for } I < 0 \tag{3.13}$$

Figure 3.7b shows the values of these equations across different values of I and V, computed using $R_{ic} = R_{id} = 0.002$.

**Voltage function:** $M$

Determining the $M$ function relies on the voltage curves found in a spec-sheet.

Let $x_1^I$ and $x_2^I$ be the initial and final charge content of the battery that is being charged or discharged at current $I$, i.e., when the voltage reaches $V_{min}$ or $V_{max}$ respectively. For

example, the $x_1^I$ and $x_2^I$ values have been labeled in Fig. 3.8 for the C/2 curve ($I = 15$ Amperes). Correspondingly, let $V^I(x)$ be the voltage when the battery charge content is $x$ while current $I$ is applied. Let $E_d^I(Q)$ be the energy drawn from a full battery when $Q$ Coulombs are discharged at current $I$, and $E_c^I(Q)$ be the energy content of the battery when it is charged to $Q$ Coulombs using current $I$. Finally, let $V^I(Q)$ be the battery voltage when current $I$ is applied and the battery reversible capacity is $Q$ Coulombs. Then,

$$E_d^I(Q) = \int_{x_1^I}^{Q} V^I(x)\eta_d(I,V)\ dx \tag{3.14}$$

$$E_c^I(Q) = \int_{x_1^I}^{Q} V^I(x)\eta_c(I,V)\ dx. \tag{3.15}$$

Subsequently, each point in the $M$ function can be calculated as:

$$M(I, E_d^I(Q)) = V^I(Q) \text{ for } I < 0 \tag{3.16}$$

$$M(I, E_c^I(Q)) = V^I(Q) \text{ for } I > 0. \tag{3.17}$$

In practice, we numerically evaluate the integrals in Eqs. (3.14) and (3.15) as Riemann sums, where the $(x, V^I(x))$ tuples are obtained by digitizing the curves from the spec sheet using a standard plot digitizer[7] and linearly interpolating between the C-rate curves using a suitably fine digitization grid. Fig. 3.5 shows the $M$ mesh for the example LTO cell. This maps the energy content in the domain [0 Wh,72.5 Wh] and current in the domain [-5 C, 5 C] to a voltage in the range [1.5 V, 2.7 V].

**Energy Content Limits:** $a_1(I), a_2(I)$

The upper and lower limits on energy content can also be derived from the voltage vs. reversible capacity curves in the spec sheet. By definition, $a_1(I)$ is the energy remaining in a battery when the voltage reaches $V_{min}$ while being discharged with current $I$, i.e., $a_1(I) = E_d^{max} - E_d^I(x_2^I)$, where $E_d^{max}$ is the maximum[8] $E_d^I(x_2^I)$ over all $I$. Similarly, $a_2(I)$ is the energy in the battery when the voltage reaches $V_{max}$ at current $I$, i.e., $a_2(I) = E_c^I(x_2^I)$. Summarizing:

$$a_1(I) = E_d^{max} - E_d^I(x_2^I) \tag{3.18}$$

$$a_2(I) = E_c^I(x_2^I) \tag{3.19}$$

[7] We use the online tool available at https://automeris.io/WebPlotDigitizer/

[8] Typically, the maximum amount of energy is obtained by using a very low discharging current

Figure 3.7a shows the energy limit functions for an LTO cell.

#### Charging/Discharging limits: $\alpha_c$, $\alpha_d$

The maximum charging and discharging C-rates may be obtained from the spec and used directly without any adjustments. $\alpha_c$ is the maximum charge current, and $\alpha_d$ is the maximum (negative) discharge current.

### 3.3.3 Evaluation

The evaluation of the PI model consists of two steps. Firstly, we validate it against the voltage curves in the spec sheet, to evaluate the degree to which the model can reproduce the spec voltage curves that it was calibrated with, effectively closing the loop between model calibration input and model output. Secondly, we validate it against an extensive dataset of battery measurements that were obtained in collaboration with the Technology Center for Energy in Ruhstorf, Germany. In both steps, we compare the PI model's accuracy to the accuracy of Tremblay's Li-ion battery model (TD model) [165].

We compare the PI model against the TD model for two reasons. First, the TD model is extensively used for battery simulation[9], and is the default battery model in Matlab's Simulink package. Second, just like the PI model, the TD model also restricts the data used for model calibration to what is available in a battery spec sheet.

Our results are computed on a laptop computer with a 2.6 GHz processor, with the model implemented in Matlab. Although optimizing the model for computational complexity was not a focus of this work, we note that thousands of model state updates have been computed within a minute.

#### Step 1: Validation Against Spec

We evaluate the ability of our model to reproduce the voltage curves found in the spec. This is done by simulating constant current charging and discharging on the PI model, using the same charging rates that were used to obtain the voltage curves from the spec and evaluate how closely the spec voltage curves are reproduced by the model.

We compare the PI model against two versions of the TD model. The first version is the default implementation of the model in Simulink. Since the default TD model does

[9] Google Scholar indicates over 500 citations to each of [165] and [166].

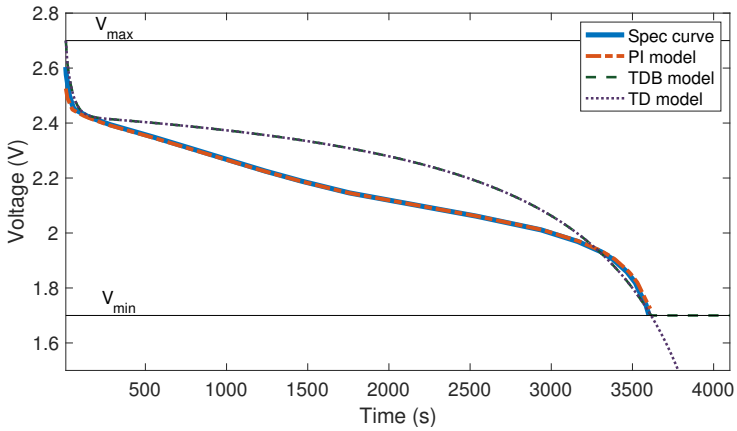


Figure 3.9: Voltage of LTO cell discharged at 1C.

not include a BMS, we created our own version of the TD model (referred to as the TDB model) that incorporates a BMS. The BMS holds the cell voltage within the battery's specified voltage range in cases where the TD model's voltage would otherwise exceed the range. This allows a fairer comparison of the TD and PI models. The simulation files and the experimental data used to evaluate our model are freely available on Github[10].

We calibrate the PI model using all the voltage curves in the spec of the LTO and $LiFePO_4$ cells; the TD and TDB models use the 1 C LTO discharge voltage curve and the 4.55 C $LiFePO_4$ discharge voltage curve[11] from the spec. Note that the LTO cell spec has voltage curves for up to 2 C discharging and 1 C charging and the $LiFePO_4$ spec has curves for up to 9.1 C discharging and no charging voltage curves.

Fig. 3.9 compares the models against the 1 C discharging curve from the LTO spec. The PI model replicates this curve very well, unlike the TD and TDB models. A similar pattern is seen in Fig. 3.10 for the LTO cell charged at 1 C. For the $LiFePO_4$ technology, Fig. 3.11 compares model performance for a 4.55 C discharging curve. All models replicate the curve well for this technology, although the TD model has higher errors at the start of the discharge.

Fig. 3.12 shows the mean absolute voltage error (MAVE) for LTO spec voltage curves. The PI model has very low error on average, especially compared to the TD and TDB models. The same comparison for $LiFePO_4$ cells is given in Table 3.3, since there are only

[10] https://github.com/iss4e/PIModel_Testing.git

[11] The voltage curve with the lowest current in the $LiFePO_4$ spec sheet was measured at 5 A, which is 4.55 C.

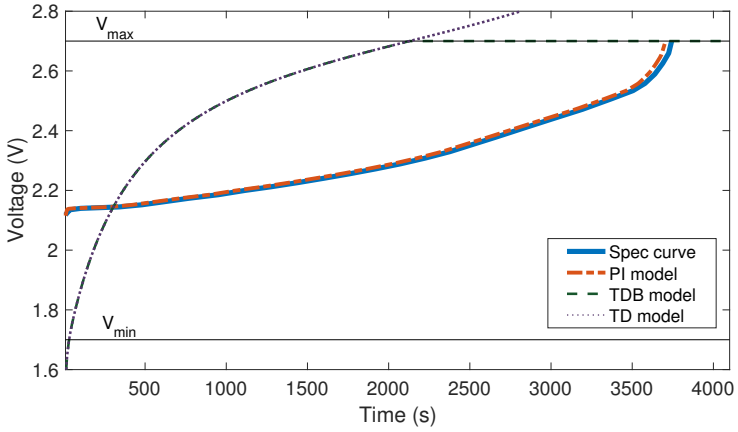


Figure 3.10: Voltage of LTO cell charged at 1 C constant current.

Table 3.3: MAVE (in Volts) for $LiFePO_4$ spec voltage curves

| C-rate | PI | TD | TDB |
|---|---|---|---|
| -9.1 | 0.012 | 0.039 | 0.039 |
| -4.55 | 0.011 | 0.12 | 0.12 |

two voltage curves to compare against in the spec of this cell. TD and TDB models behave similarly because the BMS is rarely active during these tests.

### Step 2: Validation with Experimental Data

The spec sheets for the cells used in our evaluation include reversible capacity curves for a limited number of charging and discharging currents. Furthermore, the spec sheet characterizes an average cell, but each individual cell may vary slightly depending on the manufacturing process. To better evaluate our model, we measured voltage curves over a wider range of charging and discharging currents using the setup for several cells, as described next. For example, the spec sheet for the LTO cell shows voltage curves for currents up to 2 C, but recommends a maximum charging current of 4 C; to increase the scope of our evaluation, we run experiments to obtain voltage curves for discharging rates all the way up to 5 C to test beyond the limits of the spec sheet.

We use measured voltage curves rather than curves from the spec sheet for parameter calibration of all the models being evaluated to avoid introducing errors that are caused

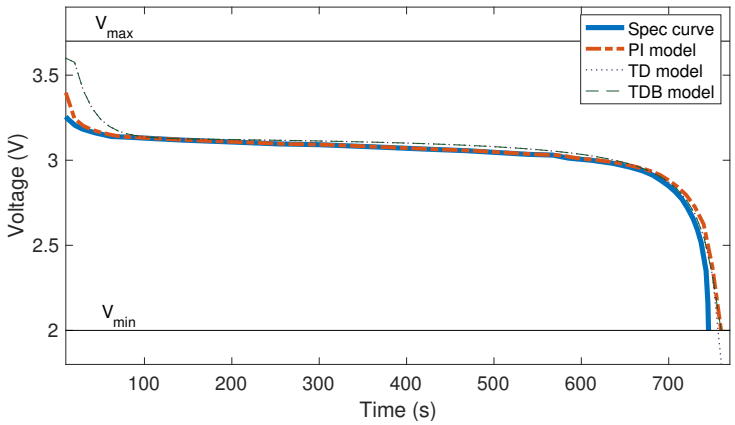


Figure 3.11: Voltage of $LiFePO_4$ cell discharged at 4.55 C constant current.

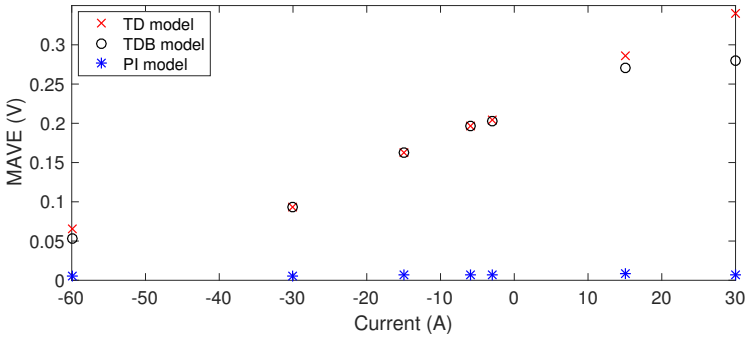


Figure 3.12: MAVE for LTO spec voltage curves.

by inconsistencies between the voltage curves of the cells in our experiments and the corresponding spec. Specifically, the PI model is parameterized using measured voltage curves across the full range of C-rates measured in our experiments, while the TD and TDB models use the 1 C discharging curve for both LTO and $LiFePO_4$ cell tests.

In our experiments with LTO cells, we run constant-current tests from 0.1 C to 4 C for both charging and discharging, as recommended by the spec, as well as to 5 C where the voltage curve is less linear and hence more complex to model. We also conduct an experiment with a 'real-world' power profile. In our experiments with $LiFePO_4$ cells, we keep the constant-current tests within 4 C charging and 10 C discharging.

(a) (b)

Figure 3.13: Experimental setup. (a) LTO cell in temperature chamber (a), and (b) two $LiFePO_4$ cells connected to the power brackets.

***Experimental Setup***

We performed a series of experiments on two LTO and two $LiFePO_4$ energy storage cells under different conditions. These experiments were designed and conducted with the assistance of technicians at the Technology Center for Energy in Ruhstorf, Germany. Each experiment consists of a single-cell battery that is charged or discharged according to an experiment profile, described next. The LTO cells [103] have a voltage range of [1.7[12],2.7] V, and nominal capacity of 30 Ah, although with low discharging rates a capacity of at least 32.7 Ah is possible. The $LiFePO_4$ cells [2] have a voltage range of [2.0,3.6] V, and nominal capacity of 1.1 Ah.

The experiments were conducted using BaSyTec XCTS Lab battery testing equipment (manufactured by BaSyTec GmbH, Germany), which has a programmable interface for specifying the charging and discharging processes of a cell, and hence can mimic a BMS programmed to prevent the battery voltage from going beyond $[V_{min}, V_{max}]$. The cells were connected to the test machine using a 4-wire interface. The equipment gives precise measurements of battery voltage and current. The cells were placed in a Binder MK 53-E2 temperature control chamber (Binder GmbH, Germany) to keep the ambient temperature at a constant 21° Celsius during testing. Fig. 3.13 shows the lab testing environment.

Most of the experiments involve charge/discharge cycling of the cell under different currents and recording the current, voltage, and total charge (Ah) every 10 seconds. We also ran experiments using a variable charge/discharge profile that reflects how a battery

[12]In the experiments with the LTO cell, we allowed the voltage to drop to 1.5 V, which is lower than the recommended minimum of 1.7 V given by the spec.

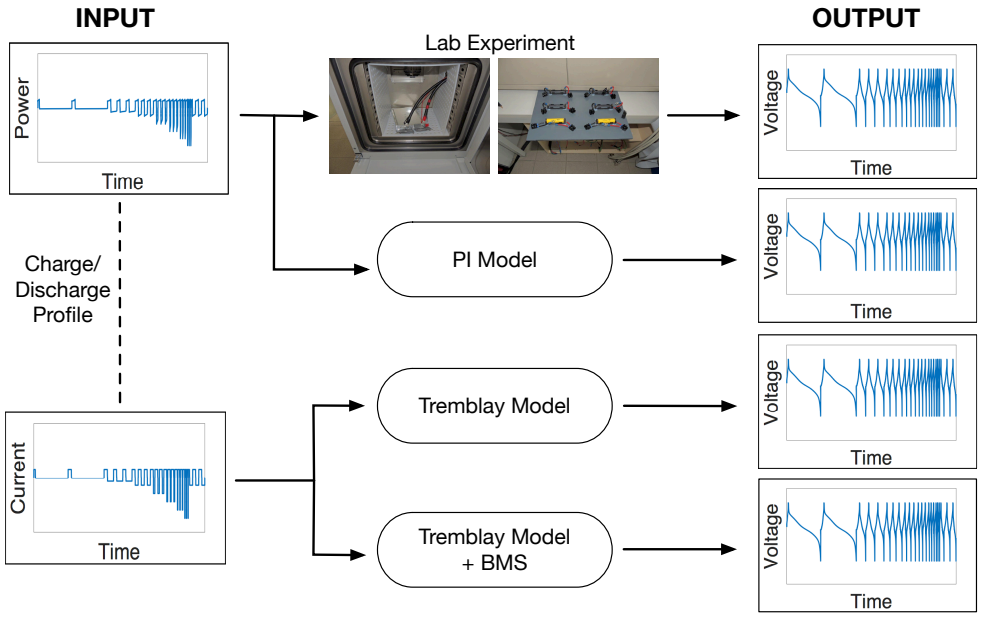


Figure 3.14: Evaluation setup. Charge/discharge profiles were tested on real cells, our PI model, and two implementations of the TD model.

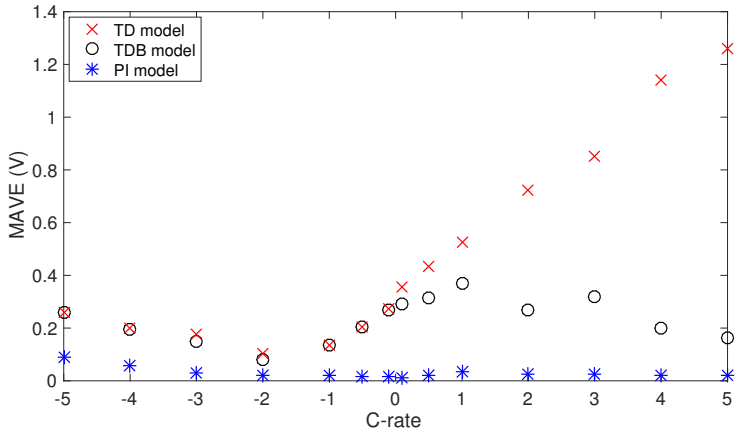


Figure 3.15: LTO MAVE with respect to measured voltage curves.

would be used to provide energy storage for a system with a solar power source and building load over over 8 hours, with a measurement granularity of 1 second. To create this profile, we took data reflecting 8 hours of solar PV power generation and a building load as measured on-site at the Technology Center for Energy at University of Applied Sciences Landshut, Germany, and used it to construct the charge/discharge pattern of a battery if it were used to buffer the excess solar energy and serve the building load. Fig. 3.14 offers a visualization of how the PI, TD, and TDB models were run on the same charge/discharge profile as the real cells in our experiments, with the measured voltage used as the benchmark for the accuracy of each model's voltage estimate.

***Constant-current tests***

The results from experiments on constant-current charging and discharging at various C-rates are summarized in Fig. 3.15 for LTO and Fig. 3.16 for $LiFePO_4$ cells. The PI model with measured parameters has a MAVE lower than 0.1 V for all C-rates tested, while the TD and TDB models have much higher errors at high C-rates compared to low C-rates.

***Real-world profile experiment***

Fig. 3.17 compares the voltage estimates of PI and TD models with the measured voltage when using a 'real-world' varying charge/discharge profile on an LTO cell. This test shows the effectiveness of the PI model beyond constant-current testing and mimics the real-world application of a battery used to balance PV generation with building power consumption. The PI model estimates the voltage with very low errors, averaging 0.016 V with spec parameters and 0.008 V with measured parameters, especially relative to the

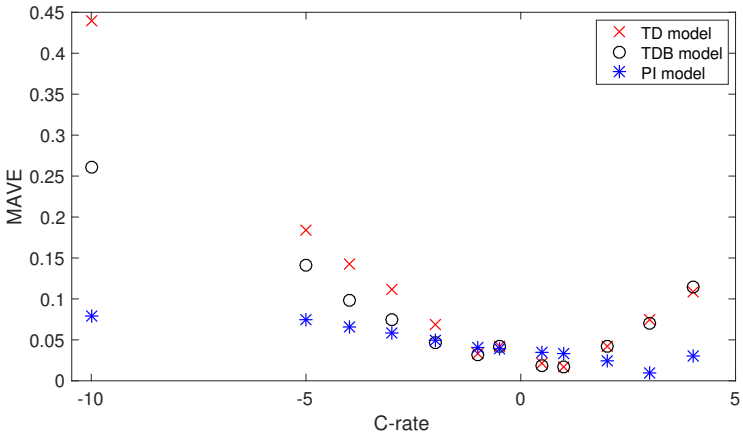


Figure 3.16: $LiFePO_4$ MAVE with respect to measured voltage curves.

errors of the TD model (0.128 V error, on average) under the same conditions. Note that the BMS is not active in this experiment, hence the TD and TDB models give the same voltage estimate.

***Review of results***

The large difference in the accuracy of the PI model compared to the TD model can be at least partially explained by noting that the TD model uses only a few points along a single discharging voltage curve for parameter calibration. While this is sufficient to characterize the voltage of the battery at that rate for some cell chemistries (as seen in Fig. 3.11), it is not sufficient for LTO cells (Fig. 3.9 and Fig. 3.10), nor is it sufficient to cover the wide range of charging and discharging rates that can be used by a real application (Figs. 3.12, 3.15, 3.16, and Table 3.3). The PI model makes use of the entire voltage curve across the full range of available charging and discharging rates, and is therefore able to accurately model voltage behaviour across the full range. Moreover, our experiments show accurate results for two Li-ion chemistries.

The PI model meets all but one of the criteria of a model well suited for system optimization studies as outlined in Section 3.2. The model is accurate, with MAVE values that are within 0.1 V for a wide range of charging and discharging rates for two Li-ion chemistries. The computation time required to compute a state update is low. The model can be calibrated using information from a spec sheet, uses power and input, and integrates the constraints of a BMS. However, the PI model is not optimization-friendly, as it has non-analytic functions for energy content limits, the voltage function, and a non-analytic process for solving the intersection between $M$ and an applied power curve.

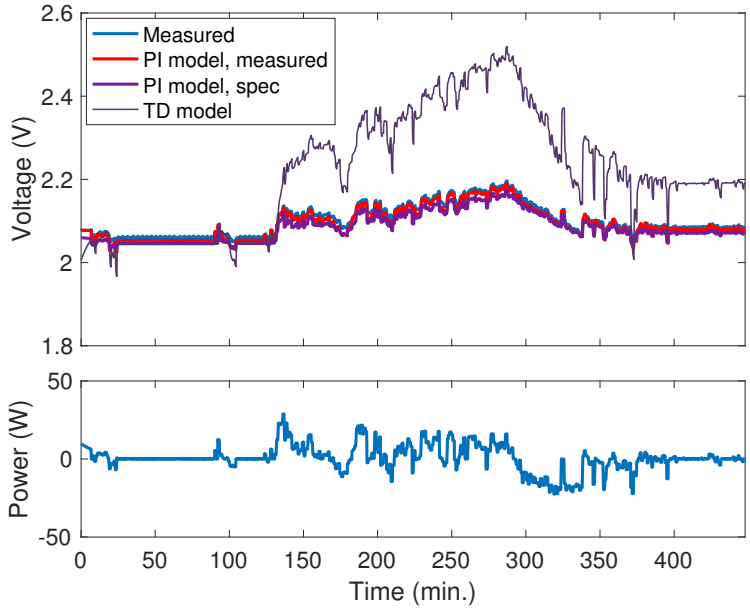


Figure 3.17: LTO battery voltage for the given charge/discharge profile.

### 3.3.4 Implementation

The PI model has been implemented as a Matlab system block that is compatible with Simulink simulation software[13]. In addition to implementing the equations of the model, the following user-friendly features have also been included.

- **Automated Parameter Computation** The parameter computation has been fully automated. The model's user specifies the files containing voltage curves as <C-rate, ampere-hour, voltage> tuples which can be extracted from the reversible capacity voltage curves found in the spec sheet using freely available online graph digitizing tools. The user also specifies the nominal capacity, internal impedance, maximum charging and discharging current, initial energy content, and the length of time steps for the discrete simulation.

[13]Matlab File Exchange link:
https://www.mathworks.com/matlabcentral/fileexchange/63078-lithium-ion-pi-model.
As of August 2019, the model has been downloaded over 500 times.

- **Smart BMS** The constraints on energy content and applied power are implemented so that the model throws a software exception when a constraint is violated. The user can adjust the power according to the information provided by the exception. The exception provides the user with the maximum power that the model could charge/discharge without violating the constraint. This allows the user to implement a 'smart' BMS that makes smooth adjustments to the input power to make the battery operation safer; we provide an example of how to capture and respond to the exception, along with the model itself, on Matlab File Exchange.
- **Parameter Extrapolation** In cases where the spec sheet fails to provide the voltage curves for the full operating range of the cell, the model uses a linear extrapolation of the available data to estimate the full range of voltage curves.

## 3.4 Models for Optimization

The PI model has two non-analytic components: the $M$ mesh function, and the energy limit functions $a_1$ and $a_2$. Both components are described as an interpolation between a set of empirically-derived points. Furthermore, the method to compute a voltage estimate using the $M$ function is also non-analytic, and cannot be expressed as a constraint to an optimizer.

In this section, using the PI model as a starting point, we replace its non-analytic components with analytic components. This process of modifying the PI model results in the development of new models that are optimization-friendly.

Our approach is to use polynomial approximations to replace the non-analytic components and to simplify non-linearities in the PI model. This task is non-trivial. The degree of the polynomials is a critical factor because higher-degree polynomials are more computationally expensive to optimize over. Non-linear – and especially non-convex – polynomial approximations would result in a complex model. However, the behaviour of a Li-ion battery is inherently non-linear. For this reason, the non-linearity of the model is hard to avoid. Therefore, in order to achieve the optimal trade-off between model accuracy and complexity, the approximations must be chosen with extreme care, keeping non-linearities which are crucial for acceptable model accuracy, discarding non-linearities which are not, and wherever possible keeping the approximations convex.

Different approximations give rise to different models. To evaluate their accuracy, in the context of their ability to estimate the energy content of the battery, we compare their results against the results of the PI model. Furthermore, we attempt to interpret what

the differences in model accuracy may mean in the context of practical applications of the models. To determine this, we conduct two system analysis case studies with the PI model, basic model, and the new models developed in this section, and compare the conclusions of the case studies.

### 3.4.1 Method of Approximation

The PI model has three principal non-trivial components: the $M$ function that maps from energy content and charge/discharge current to battery voltage, the energy content limits ($a_1(.)$ and $a_2(.)$), and the inefficiency functions ($\eta_c(.)$ and $\eta_d(.)$). The $M$, $a_1$, and $a_2$ functions are numerically derived from the laboratory experiments described in the spec sheet, i.e., they are not explicit, and this prevents the PI model from being used as part of a mathematical optimization program. The inefficiency functions are explicit but are non-linear and this increases the difficulty of using the model for mathematical optimization.

We derive a set of tractable models by approximating the three aforementioned components with polynomials. Our method is to test different combinations of approximations, starting with higher degree polynomials and working down to constant approximations, keeping only those combinations that show a noticeable difference in performance compared to the next iteration, and discarding those that do not have a favourable accuracy-complexity trade-off. We test each combination by simulating charge/discharge cycles at different C-rates and comparing the energy content estimate of the model with that of the PI model to quantify the loss of accuracy due to the approximations. The complexity of each model is on the spectrum between the PI model and the basic model; indeed, approximating the three complex components of the PI model with constants results in the basic model, from which the PI model was originally developed. In developing our models, we experiment with polynomial approximations of degree less than or equal to 3.

In this section, we show how to derive four models from the PI model, three of which are novel tractable models and have a favourable accuracy-complexity trade-off.

We now discuss our analytic approximations to voltage, inefficiency, and energy content limit functions; these are functions of the current, which is estimated from the input power using Eq. 3.9. We compute the parameters to each approximation to minimize the error between the PI model's version of the corresponding function and the approximation for a given operating range (OR), which corresponds to the maximum range of permissible currents as defined by the spec sheet or programmed into the BMS by the system operator (we elaborate on the OR in Section 3.4.2). This approach minimizes the error introduced

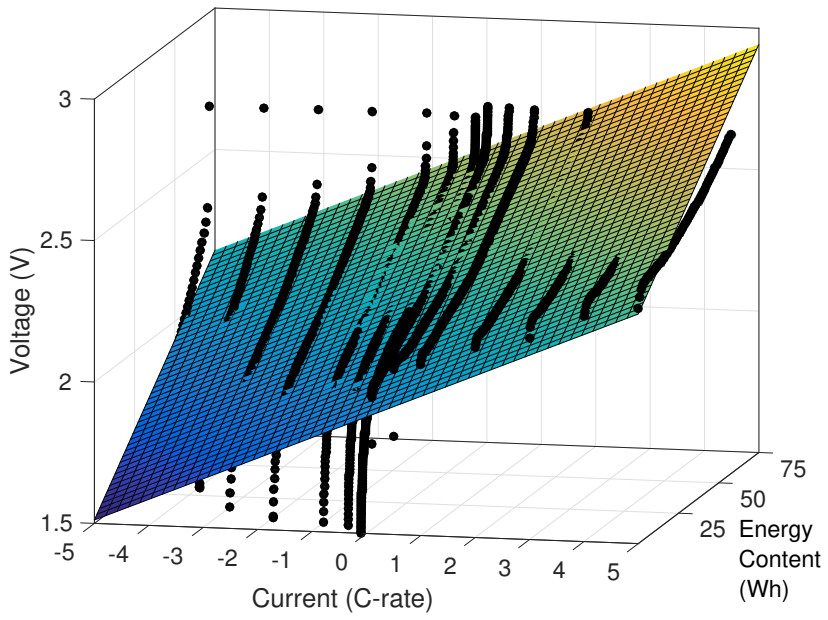


Figure 3.18: $M$ function, represented by the black dots which map the applied current and energy content to a unique voltage and are obtained from a spec sheet. The plane is a linear approximation to the $M$ function (Eq. (3.21)).

by the approximation. Each approximation is evaluated by comparing the energy content estimates of the resulting model to the PI model over constant-current charge/discharge cycles as well as a varying charge/discharge profile. We give a detailed account of the evaluation in Section 3.4.4.

### Voltage Estimate

In the PI model, the voltage for a particular energy content and applied power is estimated by using the $M$ function, which is obtained by sampling the voltage curves found in the battery's spec sheet. To make the PI model analytical, $M$ has to be expressed in analytic form. We consider the following two polynomial approximations[14]:

[14]To scale these equations to more than one cell, we divide $I(k)$ and $E(k)$ by $n_c$ wherever they appear.

$$V(k) = \begin{cases} V_{nom,d} & : P(k) < 0 \\ V_{nom,c} & : P(k) \geq 0 \end{cases} \tag{3.20}$$

$$V(k) = x_{00} + x_{10} I(k) + x_{01} E(k) \tag{3.21}$$

The constant approximation of the voltage corresponds to the 'nominal' voltage of the battery, $V_{nom,c}$ when charging and $V_{nom,d}$ when discharging; these can be computed by taking the average of the voltage curves over the OR. We note that it is possible to separate the charging voltage from the discharging voltage, as done in Eq. (3.20), for the linear approximation (Eq. (3.21)) without increasing the complexity of the model because of the natural separation between charging and discharging processes in Eqs. (3.7), (3.10), and (3.11). However, we have found that this has a negligible effect on the overall accuracy of the resulting model when applied to Eq. (3.21).

Figure 3.18 shows the non-analytic $M$ function from the PI model, and the linear approximation given by Eq. (3.21) which takes the form of a plane with the best least-squares approximation of the function[15] over the OR. The plane captures the positive slope of the voltage with increasing current and energy content but does not capture the rapid drop as the battery voltage approaches the lower limit. A cubic function can capture these dynamics. However, we found that the additional complexity of the cubic function over the plane offered negligible improvements to the accuracy of the model's energy content estimates.

### Energy content limits

Energy content limits are functions of the current, and can be obtained empirically from the voltage curves in the spec sheet. We consider the following approximations to $a_1(.)$ (similarly for $a_2(.)$):

$$a_1(I(k)) = n\bar{a_1} \tag{3.22}$$

$$a_1(I(k)) = u_1 I(k) + v_1 n \tag{3.23}$$

$\bar{a_1}$ is computed by taking the average of the $a_1(.)$ curve over the OR, and $u_1$ and $v_1$ are parameters which define the line of best fit (least-squares) to $a_1(.)$ over the OR. In our evaluation, we have found that the linear form in Eq. (3.23) offers the best trade-off between accuracy and complexity, i.e., better approximations offered very small improvement in accuracy at the cost of much higher complexity. Figure 3.19 shows the $a_1(.)$ and $a_2(.)$ functions for an LTO cell, as well as their linear least-squares and constant approximations.

[15] We use Matlab's *fit* function to compute the linear approximation.

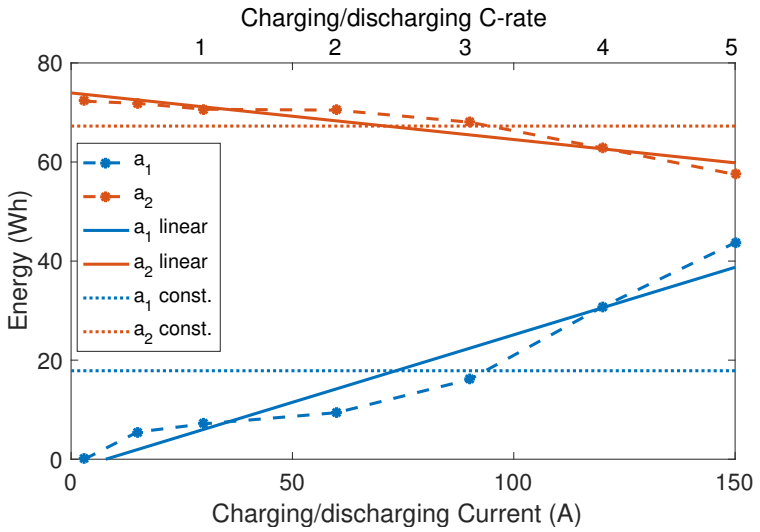


Figure 3.19: Lower ($a_1$) and upper ($a_2$) energy limits for an LTO cell, along with linear and constant approximations computed over the maximum OR of the battery (up to 5 C charging and discharging)

### Charging/discharging efficiencies

In the PI model, the charging (resp. discharging) efficiency (Eqs. (3.12), (3.13)) is a function of the voltage, current, and nominal internal resistance during charging ($R_{ic}$) (resp. discharging ($R_{id}$)). The inefficiency function makes the model inherently non-linear in two ways. Firstly, both $I$ and $V$ are variables of the model and hence the equation is quadratic in terms of the variables. Secondly, the energy content update equation (Eq. (3.7)) has a product of the efficiency and the applied power, which is the free variable in an optimization problem; this causes Eq. (3.7) to also be quadratic in terms of the variables of an optimization problem. To remove or lessen the degree of these non-linearities, we consider the following approximations for $\eta_c$ (similarly for $\eta_d$):

$$\eta_c = \bar{\eta_c} \tag{3.24}$$

$$\eta_c(I(k)) = 1 - \frac{I(k)R_i}{V_{nom,c}} \tag{3.25}$$

Using constant approximations for the charging ($\bar{\eta_c}$) and discharging ($\bar{\eta_d}$) efficiency linearizes this part of the model; they are calculated as the average values of $\eta_c(I, V)$ and $\eta_d(I, V)$ over the OR. The second approximation (Eq. (3.25)) can be obtained by using

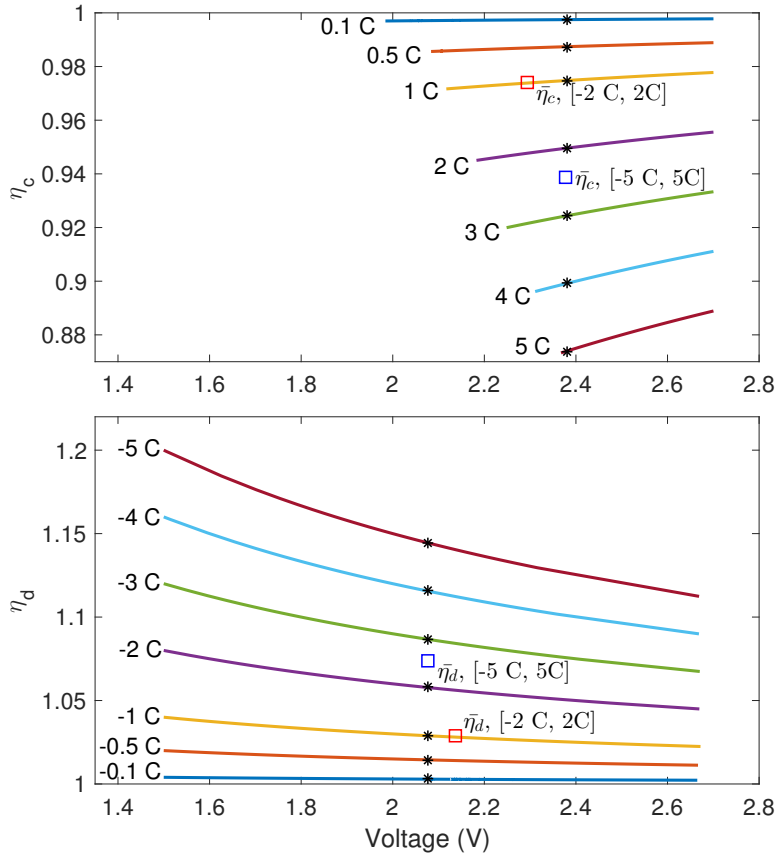


Figure 3.20: The charging and discharging efficiencies of an LTO cell with $R_{ic} = R_{id} = 0.002\ m\Omega$. Each solid line represents the efficiency calculated using Eqs. (3.12) or (3.13) at the labeled charging/discharging C-rate. Each asterisk represents the approximation where $V = V_{nom,c}$ for charging and $V_{nom,d}$ for discharging. The constant approximations $\bar{\eta}_c$ and $\bar{\eta}_d$ are shown for two labeled operating ranges.

$V_{nom,c}$ in place of $V(k)$ to reduce the dimensionality of the calculation, thereby reducing the non-linearities in the model.

To illustrate the efficiency functions and the effect of the approximations, the charging and discharging efficiencies of an LTO cell are shown in Figure 3.20, with solid lines representing Eqs. (3.12) and (3.13), asterisks representing the efficiencies when calculated using a constant approximation of the charging voltage as $V_{nom,c}$ and $V_{nom,d}$, and squares representing constant efficiency approximations.

### 3.4.2 Operating Range

The operating range (OR) is a critical factor in calibrating the parameters to the approximations of different curves. We define the OR as the range of currents expected to be applied to the battery, and "full OR" to refer to the maximum discharging and charging currents that the battery is capable of ($\alpha_d$, $\alpha_c$). This is a practical (as opposed to theoretical) consideration, since the range of currents depends on the battery application.

Accordingly, we calibrate the parameters of the model to maximize accuracy in the OR, as opposed to in the full OR. It is worthwhile to optimize the accuracy of the model in the range where it matters, even at the cost of lower accuracy in remaining parts of the full OR where accuracy does not matter. Indeed, many applications do not use the full OR. The narrower the OR – the lower the modelling error.

In Figures 3.18, 3.19, and 3.20, the approximations to the corresponding functions are computed over the full OR ([-5 C, 5 C] for the LTO cell) unless otherwise noted. In practice, we compute the constant approximations to be equal to the average value of their corresponding curve over the OR, and choose linear approximations that minimize the least-squared error to the curve over the OR. For example, Figure 3.20 shows the constant approximation for charging efficiency ($\bar{\eta}_c$) and discharging efficiency ($\bar{\eta}_d$) for the full OR of [-5 C, 5 C] as well as for a narrower OR of [-2 C, 2 C]; $\bar{\eta}_c$ is higher for the narrower OR because the efficiency is higher at lower C-rates.

### 3.4.3 Models

We have developed notation to refer to each of the four models described next, based on how the voltage, energy limits, and efficiencies are approximated. Similar to Kendall's notation for queues [83], each component of the model is described by a letter with slashes in between, i.e., Voltage function/Energy limit/Efficiency (V/E/$\eta$) notation. A letter 'Q'

Table 3.4: Summary of Analytical Models

| **Model** | | **Approximations** | | |
|---|---|---|---|---|
| | | *Voltage* | *Energy Content Limits* | *Efficiency* |
| C/C/C | $V = \begin{cases} \boldsymbol{V_{nom,d}} & : P < 0 \\ \boldsymbol{V_{nom,c}} & : P \geq 0 \end{cases}$ | | $\boldsymbol{\bar{a}_1} n$<br>$\boldsymbol{\bar{a}_2} n$ | $\boldsymbol{\bar{\eta}_d}$<br>$\boldsymbol{\bar{\eta}_c}$ |
| C/L/C | $V = \begin{cases} \boldsymbol{V_{nom,d}} & : P < 0 \\ \boldsymbol{V_{nom,c}} & : P \geq 0 \end{cases}$ | | $a_1 = \boldsymbol{u_1} \dfrac{P}{\boldsymbol{V_{nom,d}}} + \boldsymbol{v_1} n$<br>$a_2 = \boldsymbol{u_2} \dfrac{P}{\boldsymbol{V_{nom,c}}} + \boldsymbol{v_2} n$ | $\boldsymbol{\bar{\eta}_d}$<br>$\boldsymbol{\bar{\eta}_c}$ |
| C/L/L | $V = \begin{cases} \boldsymbol{V_{nom,d}} & : P < 0 \\ \boldsymbol{V_{nom,c}} & : P \geq 0 \end{cases}$ | | $a_1 = \boldsymbol{u_1} \dfrac{P}{\boldsymbol{V_{nom,d}}} + \boldsymbol{v_1} n$<br>$a_2 = \boldsymbol{u_2} \dfrac{P}{\boldsymbol{V_{nom,c}}} + \boldsymbol{v_2} n$ | $\eta_d = 1 - \dfrac{PR_{id}}{n\boldsymbol{V_{nom,d}^2}}$<br>$\eta_c = 1 - \dfrac{PR_{ic}}{n\boldsymbol{V_{nom,c}^2}}$ |
| L/L/Q | $V = \boldsymbol{x_{00}} + \dfrac{\boldsymbol{x_{10}}}{n} I + \dfrac{\boldsymbol{x_{01}}}{n} b$ | | $a_1 = \boldsymbol{u_1} I + \boldsymbol{v_1} n$<br>$a_2 = \boldsymbol{u_2} I + \boldsymbol{v_2} n$ | $\eta_d = 1 - \dfrac{IR_{id}}{n\boldsymbol{V}}$<br>$\eta_c = 1 - \dfrac{IR_{ic}}{n\boldsymbol{V}}$ |

Note: $x_{00}$, $x_{01}$, $x_{10}$, $u_1$, $u_2$, $v_1$, and $v_2$ are fitted. **Bolded** parameters are affected by the OR.

means that that calculation uses a quadratic approximation, such as the efficiency function in the PI model; 'L' represents a linear approximation, and 'C' a constant approximation[16].

The four analytic models are summarized in Table 3.4, and described below. Three out of four models approximate the voltage using two constants $V_{nom,c}$ and $V_{nom,d}$. For these three models, we can write their formulation in terms of power by replacing $I(k) > 0$ by $P(k)/V_{nom,c}$ and $I(k) < 0$ by $P(k)/V_{nom,d}$ in the PI model and can remove Eqs. (3.8) and (3.9) completely.

[16]Technically, the voltage approximation given in Eq. (3.20) corresponds to two constants instead of one.

### C/C/C

This model is the basic model summarized in Eqs. 3.2-3.5, and we henceforth refer to it as the C/C/C model. It uses constant approximations for efficiencies and energy limits, and a constant approximation of the voltage. It is linear (see the discussion in Section 3.2.5) and all its equations can be written in terms of $E(k)$ and $P(k)$ only. It has low energy estimation errors only when the OR is narrow (as we show in [81]), which increases the importance of calibrating it with respect to the OR.

Treating the basic model as a derivative of the PI model makes the three-step calibration procedure clear:

1. Compute the parameters of the PI model from the spec.
2. Compute the constant approximations of the PI model's parameters.
3. Use the constant approximations as the values for the corresponding parameters of the basic model.

Since the PI model's calibration uses all the available spec data, and since using this data is an important factor for model accuracy, this calibration method improves the accuracy of the basic model. For example, the charging efficiency calculated for an LTO battery with an OR of [-5 C, 5 C] is 0.94, which increases to ≈0.975 for an OR of [-2 C, 2 C]. For a $LiFePO_4$ battery [2], the efficiency for an OR of [-2 C, 2 C] is calculated to be ≈0.985.

### C/L/C

This model uses constant approximations of the charging and discharging voltage as well as charging and discharging efficiencies, and uses a linear approximation to the energy content limits. Similarly to the C/C/C model, all the equations can be written in terms of $E(k)$ and $P(k)$ only. Notably, it has the same complexity as the C/C/C model (linear), while offering a better approximation to the energy limits.

**C/L/C:**

$$E(k) = E(k-1) + \Delta_E(k) \tag{3.26}$$

$$\Delta_E(k) = \begin{cases} \bar{\eta_c} P(k) T_u & : P(k) \geq 0 \\ \bar{\eta_d} P(k) T_u & : P(k) < 0 \end{cases} \tag{3.27}$$

$$n\alpha_d V_{nom,d} \leq P(k) \leq n\alpha_c V_{nom,c} \tag{3.28}$$

$$u_1 \frac{P(k)}{V_{nom,d}} + v_1 n \leq E(k) \leq u_2 \frac{P(k)}{V_{nom,c}} + v_2 n \tag{3.29}$$

### C/L/L

This model uses constant voltage approximations, the same inefficiency functions as the PI model, and linear energy content limits. The constant voltage approximation makes the efficiency function linear. In the formulation, we have excluded an explicit $I(k)$ term since, due to the constant voltage approximations, the current is estimated as either $P(k)/V_{nom,c}$ or $P(k)/V_{nom,d}$. Hence, all the equations can be written in terms of $E(k)$ and $P(k)$. Altogether, this model is quadratic in $P(k)$.

**C/L/L:**

$$E(k) = E(k-1) + \Delta_E(k) \tag{3.30}$$

$$\Delta_E(k) = \begin{cases} \eta_c(P(k)) P(k) T_u & : P(k) \geq 0 \\ \eta_d(P(k)) P(k) T_u & : P(k) < 0 \end{cases} \tag{3.31}$$

$$\eta_c(P(k)) = 1 - \frac{P(k) R_{ic}}{n V_{nom,c}^2} \tag{3.32}$$

$$\eta_d(P(k)) = 1 - \frac{P(k) R_{id}}{n V_{nom,d}^2} \tag{3.33}$$

$$n\alpha_d V_{nom,d} \leq P(k) \leq n\alpha_c V_{nom,c} \tag{3.34}$$

$$u_1 \frac{P(k)}{V_{nom,d}} + v_1 n \leq E(k) \leq u_2 \frac{P(k)}{V_{nom,c}} + v_2 n \tag{3.35}$$

### L/L/Q

This is the most complex of our explicit models. It combines a linear voltage function, the same quadratic efficiency formula used in the PI model, and a linear approximation for

the energy limits. The formulation resembles that of the PI model with the voltage and energy limits replaced by analytic functions.

**L/L/Q:**

$$E(k) = E(k-1) + \Delta_E(k) \tag{3.36}$$

$$\Delta_E(k) = \begin{cases} \eta_c(P(k))P(k)T_u & : P(k) \geq 0 \\ \eta_d(P(k))P(k)T_u & : P(k) < 0 \end{cases} \tag{3.37}$$

$$\eta_c(I(k)) = 1 - \frac{I(k)R_{ic}}{nV(k)} \tag{3.38}$$

$$\eta_d(I(k)) = 1 - \frac{I(k)R_{id}}{nV(k)} \tag{3.39}$$

$$V(k) = x_{00} + \frac{x_{10}}{n}I(k) + \frac{x_{01}}{n}E(k) \tag{3.40}$$

$$I(k) = \frac{P(k)}{V(k)} \tag{3.41}$$

$$n\alpha_d \leq I(k) \leq n\alpha_c \tag{3.42}$$

$$u_1 I(k) + v_1 n \leq E(k) \leq u_2 I(k) + v_2 n \tag{3.43}$$

In this model, all the equations can be written in terms of $E(k)$, $I(k)$ and $P(k)$. The model is non-linear in both $I(k)$ and $P(k)$, which makes this model much more complex and difficult to use.

### 3.4.4 Evaluation

For each of the analytic models discussed in Section 3.4.3, we quantify the loss of accuracy in the energy content estimate with respect to the PI model because, as is shown in Section 3.3, it accurately reflects battery dynamics.

To evaluate our models, we use LTO [103] and LiFePO4 [2] battery measurements collected while running charge/discharge cycling experiments on these batteries at different C-rates with 2 or 3 cycles per C-rate, and use the PI model to calculate the energy content of these batteries over the course of the experiment. Next, we recreate these experiments using simulations of each one of our derived analytic models, assuming the same applied power that has been used in our experiments. Finally, we compare the energy content in the battery over time as simulated by each of our models to the energy content calculated by simulating the battery using the PI model.

### Constant-current cycles

For this section of the evaluation, we calibrate our models over the full operating range of the batteries: [-10 C, 4 C] for $LiFePO_4$, and [-5 C, 5 C] for LTO. Figure 3.21 shows the mean absolute energy error (MAEE) over a full charge (or discharge, for negative C-rates) as a percentage of the maximum battery capacity for each of our analytic models with respect to the PI model. The order of the models from most accurate to least accurate mimics their order from most complex to least complex; in other words – the more complex the model, the more accurate it is.

The relative differences in accuracy between the models can be used to infer the contribution of each approximation to the modelling error. For example, a linear voltage function (L/L/Q) offers little improvement in accuracy over using two constants (C/L/L) for all but the most extreme C-rates.

C/C/C and C/L/C models have their lowest errors in the 2-3 C range for LTO, for both charging and discharging, and at 5 C discharging and 2 C charging for LiFePO4. We infer that, as is shown for an LTO cell in Figure 3.20, the constant efficiency approximation is close to the PI model's efficiency estimate at these currents.

It is important to note that errors introduced by two approximations may compensate each-other. For example, one approximation which causes the model to over-estimate the energy content may be compensated by another approximation which causes the model to under-estimate the energy content. In our evaluation, this appears to be the case in the 5 C charging test for the LTO battery, where the errors of the C/C/C and C/L/C models are unexpectedly lower than the C/L/L model error.

### Varying charge/discharge profile

To evaluate the accuracy of the models in realistic conditions, we use data from the experiment which mimics real battery behaviour in a system with PV panels and building load, as described in Section 3.3.3. We simulate this system using our analytic models which have been calibrated with an OR of [-2 C, 2 C], and with the PI model as our benchmark. Figure 3.22 shows the power profile used in our experiment, the resulting energy content evolution over an 8-hour period, as well as the residual with respect to the PI model's energy content estimate. This experiment highlights the effect of the efficiency and voltage function approximations, since the battery did not approach any energy content limits over the course of the test. The residual for all the models stays below 1% of the total energy capacity of the LTO battery, with the residual of C/L/L and L/L/Q models well below 0.1%.

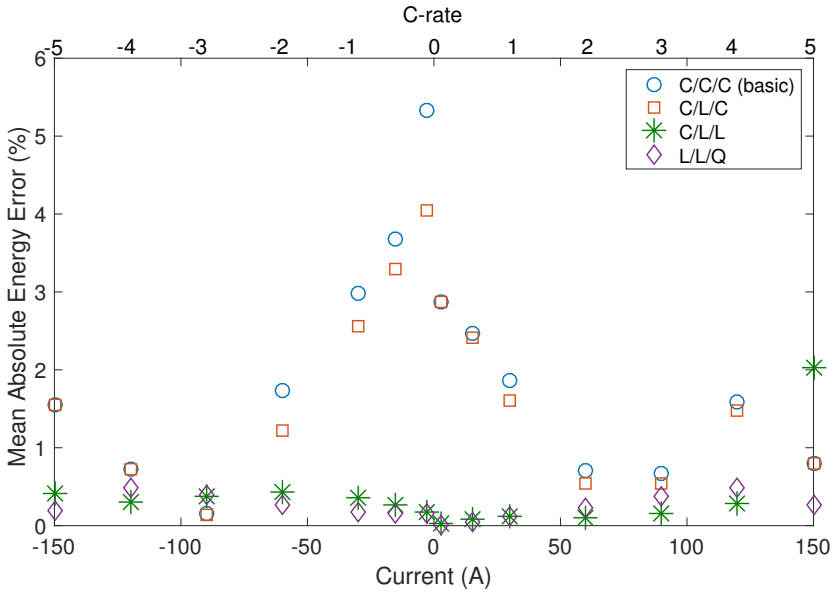


(a) LTO

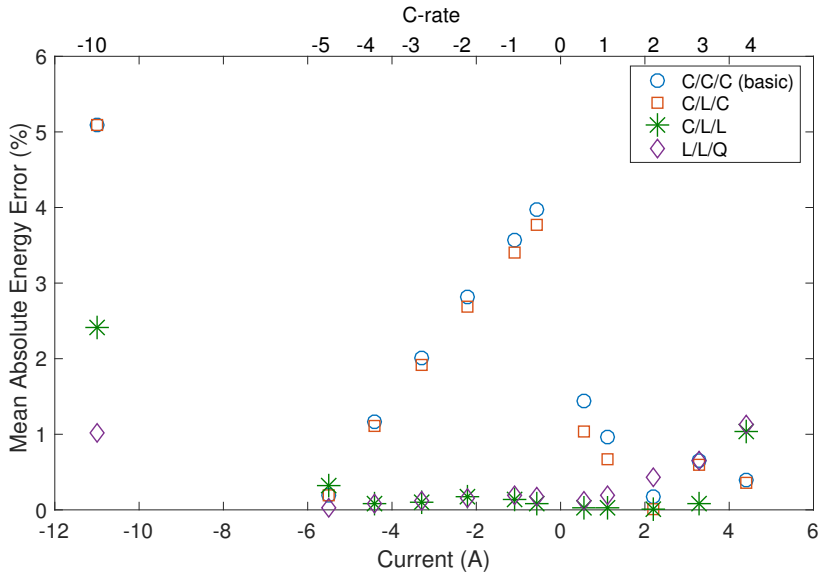


(b) $LiFePO_4$

Figure 3.21: Mean absolute energy error for LTO cell (3.21a) and LiFePO4 cell (3.21b) for constant-current cycles at different C-rates (top scale)/Amperes (bottom scale).

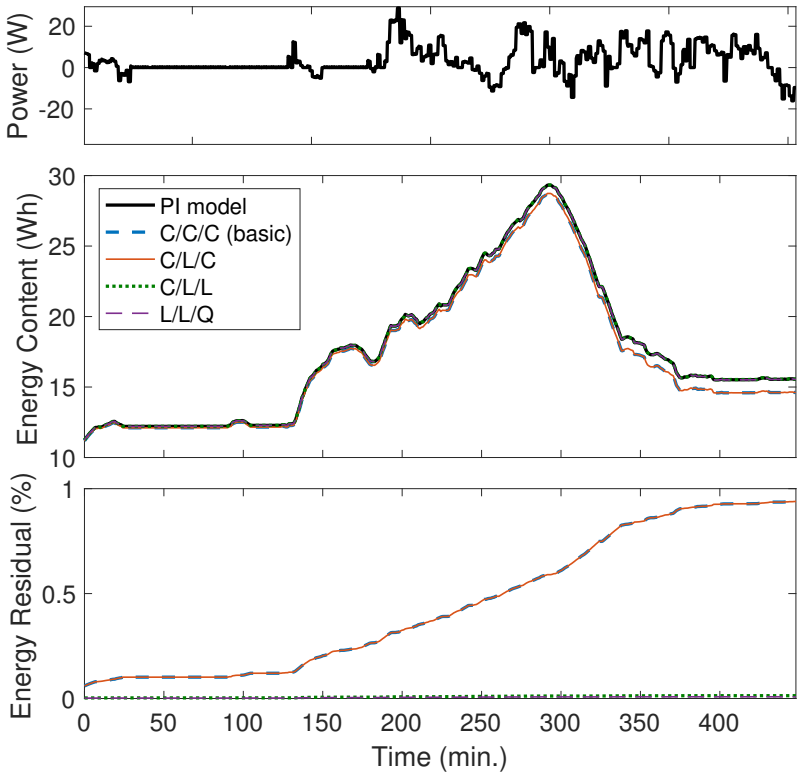


Figure 3.22: The energy content over 8 hours in a realistic deployment charge/discharge profile. The residual between the analytic models and the PI model (bottom), as well as the applied power trace (top), are also shown.

From this experiment alone, it is unclear whether residuals of 1% or 0.1% are sufficiently low for obtaining useful results where these models are used for energy system analysis. Likewise, the impact of MAEE on the practical application of these models is not sufficiently clear. In order to resolve this problem, we compare these models in terms of their performance in two case studies.

### 3.4.5 Application Sensitivity Study

We perform case studies that compare the accuracy of our models with respect to the PI model for relevant metrics in two battery applications: a solar farm application, and a grid regulation services application. The purpose of these case studies is to exhibit the effects of the choice of battery model on the conclusions that are made. Both case studies are performed via simulations and are described below.

To obtain results with models which are calibrated using the proper OR, which differs for a given application and battery size[17], we run each simulation twice. Firstly, we calibrate the model with the full OR of the battery, i.e., $[\alpha_d, \alpha_c]$, and simulate the application. Secondly, we observe the OR that was actually used by the simulated application, re-calibrate the model, and run the simulation again. We report the results from the second simulation.[18]

#### Solar farm firming

We consider a small (100 kW) solar farm participating in an electricity market with hourly power commitments and penalties for failing to adhere to the commitment. Figure 3.23 shows the basic structure of the system, where $S(t)$ is the incoming solar power which we model using measured data. The load $L(t)$ represents the hourly power commitment made by the operator. We assume that the operator always commits to providing the average hourly average solar production, and a battery is used to firm (or to level) the intra-hourly fluctuations in PV production.[19]

[17] The effect of the battery size on the OR, which is given in terms of C-rate, can be explained concisely via example: an application may require a current of 100 Amps, which equates to 1 C for a particular battery, but only 0.5 C when the battery's size is doubled

[18] One could continue re-running simulations using the OR observed in the previous simulation until the OR used to obtain parameters and the OR observed in the simulation converge. In practice, we found that one iteration of this method is enough to get a reasonable OR.

[19] A comprehensive study of this system can be found in [48].

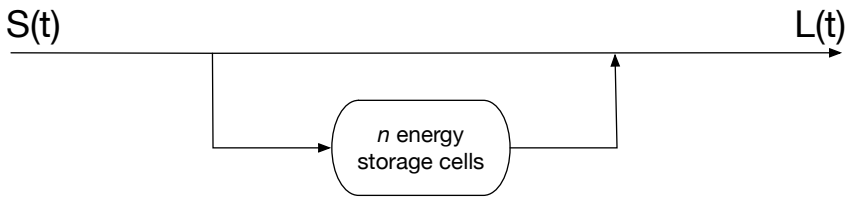


Figure 3.23: Power flow diagram in a solar farm equipped with a battery. S(t) is the power generated by the solar cells, and L(t) is the load on the system (power commitment).

An important metric in this application is the amount of energy that was needed from the battery but not delivered, i.e., when an insufficient amount of PV power is generated to meet the power commitment and the operator needs to discharge the battery but is unable to because the battery does not have enough energy or the required power is too high for the battery to provide. We call this the *unmet load* on the battery. Choosing a battery size to meet an unmet load target is a classic optimization problem. To see the difference that the choice of battery model will make on the conclusions of this study, we calculate the unmet load by simulating the system with a solar measurement trace for 100 days across a range of battery sizes.

We present results with the model calibrated for the LTO battery chemistry and remark that we have obtained similar results for the $LiFePO_4$ chemistry. Figure 3.24 shows the unmet load calculated for different battery sizes using the PI model as well as the four analytic models, and the OR used to calibrate model parameters for each battery size. The relative error is very low for the L/L/Q and C/L/L models, and very high for C/C/C and C/L/C models.

Even though the C/L/C model has a low error for 10 and 50 cell batteries, its error is high for larger sizes; this is due to the OR and its effect on the various model parameters. For smaller battery sizes, the OR is maxed out at [-5 C, 5 C], although the actual currents are much more uniformly distributed across the OR with a small (10 to 50 cell) battery compared to the 100-cell battery where the 5 C is a rare occurrence. The OR gradually narrows as the number of cells increases to 300. For a large OR, it is crucial to model the energy limits as a function of the current, which is why the C/L/C model yields more accurate results than the C/C/C model. For narrower OR, the efficiency estimate becomes a more important factor than the factor of energy limits, as demonstrated by the high accuracy of L/L/Q and C/L/L models and the low accuracy of C/L/C and C/C/C models.

Observe that the conclusions of this study can vary greatly depending on the choice

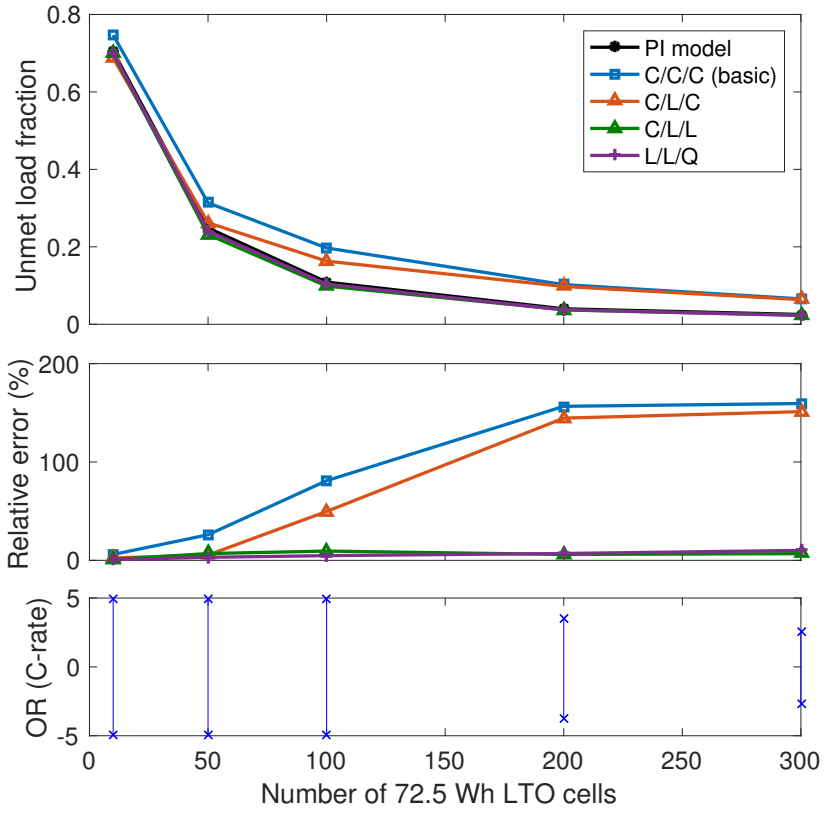


Figure 3.24: Unmet load for different battery sizes, computed using five different models. Relative error is with respect to the PI Model, and the effective OR calculated by the PI model is shown in the bottom-most figure.

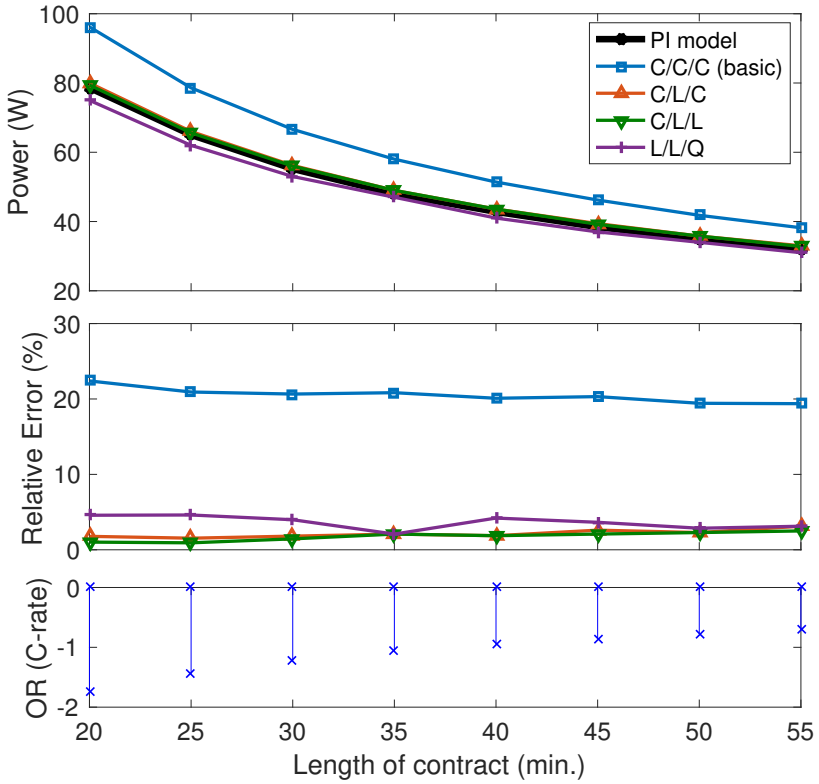


Figure 3.25: Maximum power commitment for guaranteed delivery, shown for various contract lengths. The relative error with respect to the PI model is also shown. The OR figure shows the actual OR determined from the PI model simulation.

of battery model. For example, if the purpose of the study is to determine the smallest battery size that is sufficient to achieve an unmet load target of 40%, all models suggest a similar size. However, for a more realistic 10% target, the sizes suggested by the C/C/C and C/L/C models are over twice as large as the sizes suggested by the C/L/L and L/L/Q models. We interpret this as a confirmation that neither of the linear models works well in this application, and a more complex model such as C/L/L is needed to obtain accurate results.

### Regulation

We now study the impact of choice of battery model on a grid regulation application[20] where the battery operator declares the amount of power that they can provide or absorb (by discharging or charging the battery, respectively) for the duration of a short-term contract. Given the amount of energy in the battery, the operator needs to calculate the maximum amount of power regulation they could provide in the worst case where their power commitment is requested for the entire duration of the contract. In this case study, we model an LTO battery and assume that the energy content at the start of the contract is 50% of the battery's maximum capacity. We present the results for down-regulation, where the battery is only discharged. Figure 3.25 shows the power that the battery can consistently provide for various contract durations. We also compute the relative error with respect to the results obtained using the PI model.

The C/C/C model has very poor accuracy for all contract lengths, which may be explained by noting that the energy limits are the most important model parameters for this application. The L/L/Q model has slightly greater (up to 5%) errors than the C/L/C and C/L/L models, which is not an intuitively expected result and is an important observation. The explanation for this anomalous behaviour is that, since we are dealing with approximations of multiple components of the PI model, there is room for errors introduced by these components to compensate each other. Indeed, the improved voltage estimate of the L/L/Q model in turn leads to an accurate current estimate, which is then used to calculate efficiency and energy limits, while the constant voltage approximation in C/L/L and C/L/C models leads to a poor current estimate which actually ends up partially negating the errors in the energy limit approximation.

The C/L/C model has low errors across all contract lengths. We interpret these results as a confirmation that a linear model (C/L/C) is sufficiently accurate for this application.

To highlight the effect of the OR on the accuracy of the models, we have calculated the results for the regulation application with a full OR ([-5 C, 5 C]) for all contract lengths. Figure 3.26 shows the power commitment suggested by the C/C/C, C/L/C, and C/L/L models with a full OR, compared to the same models calibrated with the actual OR of the application for each contract length. Every model has lower accuracy when calibrated for the full OR. Note that the L/L/Q model is not shown because its accuracy is almost unchanged when calibrating for a narrower OR; this is an expected result, as the L/L/Q model uses more accurate approximations which are not as sensitive to the OR.

[20] A comprehensive study of this application can be found in [43]

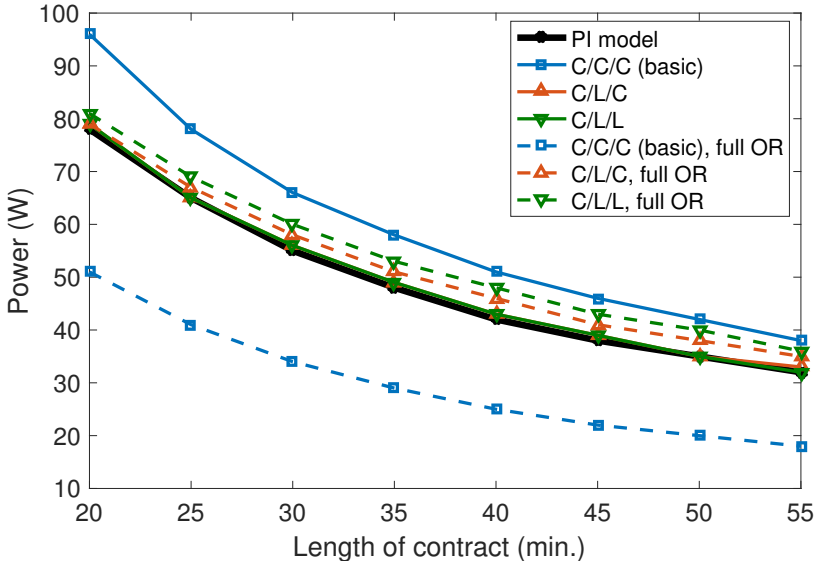


Figure 3.26: Computed vs. full OR calibration for the regulation application.

### 3.4.6 Guidelines for Model Selection

Our extensive experimentation with different models and calibration methods has given us insight into how to choose and calibrate a battery model for energy system analysis, in particular for system control and sizing optimization studies. Below, we outline four considerations that are important for selecting and calibrating a model.

1. *What is the size of the optimization problem?* For small problems – where the problem size corresponds to the number of variables – that could be solved by non-linear optimization methods in a reasonable time frame, the L/L/Q and C/L/Q models could be used. Large problems require more efficient methods such as linear programming, which in turn requires a linear model. For large optimization problems, we suggest using the C/L/C model, since it is typically more accurate than the commonly used C/C/C model, with no added complexity.

2. *What battery chemistry is being modelled?* Some chemistries, such as $LiFePO_4$, have "flat" voltage profiles which can be approximated well as a constant voltage. Cells with very low internal impedance have high efficiency values (close to 1) which could be closely approximated using a constant. The area under the voltage curve for different C-rates on a voltage vs. charge graph (commonly found in a spec sheet) can

hint at the shape of the energy limit function[21]. This information can reveal which approximations would work well, and help to guide the choice of battery model.

3. *Which constraints of the battery is the application most sensitive to?* The answer to this question may or may not be known *a priori*; if known, it can guide the selection of the model based on the effectiveness of the approximations to each model component. For example, if efficiency is a crucial parameter, we suggest using the L/L/Q or C/L/L models which have more accurate efficiency functions than the C/L/C model.

4. *What is the operating range of the application?* If the OR is known *a priori*, it should be used to calculate the parameters of the battery. If it is unknown, the analysis should first be conducted with the model calibrated using the maximum OR that the battery allows, keeping track of the OR that is actually used by the battery in the analysis. Then, the analysis should be re-run with parameters computed for the OR observed in the previous stage. This can be repeated until the calibrated OR matches the observed OR; in our experience, we have noticed that there has been little or no benefit in going beyond one or two iterations of this method.

## 3.5 Limitations and Evolution of Models

Even though the PI model and its derivatives developed and described in this chapter take into account many important factors affecting battery dynamics, they are not perfect. Understanding the limitations of these models is important for their proper use, as well as for their further development. In this section, we briefly discuss several aspects that are not taken into account in our models, and which signify possible directions of model evolution. In particular, we discuss the factors of temperature, battery degradation, and battery technologies.

### 3.5.1 Impact of Temperature

The models developed in this chapter do not take battery temperature into account. Battery temperature affects the voltage and resulting energy capacity. Cold temperatures increase the internal resistance of the battery [102], which, in accordance with Ohm's law,

[21]The area under a discharging voltage vs. charge curve is equal to the energy discharged from the battery at the given C-rate.

causes the voltage to drop and subsequently reach its lower limit quicker when discharging. The effective result is a decrease in battery energy capacity. Warmer temperatures have the opposite effect. Furthermore, charging and discharging a battery cause changes in battery temperature, since these processes are not perfectly efficient and some energy is lost as heat.

Effects of temperature in batteries have been studied in existing work, for example in [183] where an electro-thermal battery model is developed and experimentally validated; this model is calibrated via experiments to determine specific heat capacity and transfer properties of the cells and is based on differential equations. Temperature effects in a high-level model, such as the PI model and its approximations, introduce a fourth dimension into what is currently a three-dimensional inter-relationship between voltage, energy content, and charging/discharging current. These effects also introduce two additional state variables related to battery temperature and ambient temperature. The data describing the relationship between temperature, voltage, energy, and current, is not present in most of the battery spec sheets we have encountered, which increases the challenge of keeping the model spec-based and easy to use. The inclusion of temperature into our models is a possible direction for their improvement.

### 3.5.2 Impact of Battery Degradation

Battery degradation is another factor that is not considered in the models described in this chapter. The energy capacity and efficiency of a Li-ion battery gradually decrease with time and use. Factors which cause battery degradation include the loss of available $Li^+$ ions that shuttle back and forth between the anode and cathode, and the increase in internal resistance due to degradation of electrolyte. The rate of battery degradation depends not only on the time and the total amount of energy passing through the battery but also on the temperature and depth of discharge of each charge-discharge cycle.

Some detailed electrochemical models attempt to explicitly model the processes of battery degradation [129]. Even high-level models can approximate degradation either as a linear function of the total amount of energy passing through the battery [44], or by taking into account the cumulative effects of the depth of discharge of each cycle [179, 93]. Information on the expected rate of battery degradation may sometimes be found in battery spec sheets, in the form of the number of charge-discharge cycles after which the battery capacity is expected to drop to 80% of its original value, or as a curve showing the progression of the capacity fade. In both cases, degradation is rated at some nominal temperature, charging and discharging rates, and depth of discharge. The addition of time

and cumulative cycles to the parameters of the PI model is a potential direction of its evolution. These parameters could be used to set the rates of decrease in energy capacity and increase in internal resistance over the course of a system simulation.

### 3.5.3 Various Battery Technologies

The adaptation of our models to various established and emerging battery technologies, for example lead-acid, sodium-nickel-chloride, and redox-flow batteries, is possible. The PI model and its approximations have been evaluated on a set of experiments consisting of charge-discharge cycles over a wide range of C-rates for two Li-ion chemistries. Similar experiments could be conducted on other Li-ion chemistries and other battery technologies, which would serve to validate the PI model's accuracy for these technologies or form a starting point for the adaptation of model equations to match the specific characteristics of each technology. New approximations of the resulting model could then be computed using the method described in Section 3.4.1.

**Chapter 3 Summary**
Battery models are crucial for the study of PV-storage system optimization. In our work, we develop the PI model, which can be easily calibrated, has high accuracy, and is suitable for simulation studies. We then derive three analytic models which represent different trade-offs between accuracy and complexity and are suitable for optimization studies via mathematical programming. All of our models are used in two PV-storage system case studies to show the effects of each battery model on the conclusions of the studies.

# Chapter 4

# Optimization of System Sizing

In recent years, the prices of solar panels and storage have dropped dramatically, putting them in reach of many consumers. Companies such as Trina, Yingli, and Canadian Solar offer solar panels at a cost of less than USD 0.5/Watt, though installation costs may be much higher, and companies such as Tesla, Sonnen, and Moixa provide off-the-shelf (albeit still relatively expensive) storage solutions.

For an entity that would like to purchase and install solar PV panels and storage to partly or completely offset grid usage[1], the primary question is: how much of each should they buy? If the budget is not a constraint, then both can be generously sized, with ample slack capacity. However, given the high cost of system components, budget is often a binding constraint. Thus, there is a demand for practical guidance on the smallest possible sizing to adequately meet the anticipated load. By sizing, we refer to the power/energy size of the storage in kWh and the size of solar generation in kW. We expect many entities to face such a sizing problem in the future. These include individuals, small companies, and building operators faced with the rising cost of grid-provided electricity.

PV-storage systems represent a substantial monetary investment. The expected system performance, i.e., the amount of load that it will offset, is an important metric for prospective system owners. Providing a guarantee on system performance is a challenging task, given that the investment pays out over a decades-long system lifetime, and the operating environment is subject to change over time. The expected ROI is dependent on where

[1]The former case corresponds to that of an entity that remains grid-connected but wants to reduce its overall cost for electricity and the latter corresponds to an off-grid scenario. We treat them both identically in our work.

the system is deployed[2], as well as by changes in policy affecting feed-in tariff rates and electricity rates charged by utility companies, both of which are hard to predict precisely. However, an inherent feature in any deployment scenario is that the ROI is maximized when the system is sized to meet a given performance target at a minimum capital cost.

Developing methods to provide guidance on sizing is the subject of this chapter. It includes an overview of existing sizing methods in Section 4.1, research on methods for robust system sizing in Section 4.2, and a sketch of potential avenues for future developments in Section 4.3.

The specific contributions of this chapter are as follows. We develop a novel technique for robust and practical PV-storage system sizing based on system simulation and mathematical programming. We evaluate the robustness and computational complexity of these approaches, and compare them to a stochastic network calculus sizing approach, in a realistic setting. Materials in this chapter have been published in [74].

## 4.1 Related Work on Energy System Sizing

Our organization for the review of the literature on studies that optimize system components is based on two distinct aspects. The first aspect: the techniques used to optimally size the system. The second aspect: the specific deployment context of the system.

### 4.1.1 Sizing Methods

Depending on the method of optimization, the existing literature on the joint sizing of energy storage and renewable energy sources may be divided into three groups: mathematical programming, simulation, and analytical methods. We sketch these approaches here with a survey of representative work.

#### Mathematical Programming

Mathematical programming computes the optimal system sizing for a given pair of trajectories (i.e, a time series) for load and PV generation. It requires modelling the system as a set of parameters and variables that are constrained to represent the capabilities of the underlying physical system and an objective function representing the optimization target.

[2] A thorough analysis of PV-storage system ROI in different jurisdictions can be found in [76].

Table 4.1: Sizing Notation

| *Name* | *Description* |
|---|---|
| $B$ | Energy capacity of battery |
| $b$ | Number of possible battery sizes |
| $C$ | Power generation capacity of PV panels |
| $c$ | Number of possible PV panel sizes |
| $D_i$ | Load trace for scenario $i$ |
| $\delta(t)$ | Unmet load in time slot $t$ |
| $\epsilon$ | Loss of load target fraction |
| $\gamma$ | Confidence measure for loss of load probability or expected unserved energy target |
| $K_i$ | Sizing curve corresponding to scenario $i$ |
| $\lambda$ | Chebyshev bound coefficient of standard deviation |
| $\mu$ | Mean of probability distribution |
| $n$ | Number of scenarios considered |
| $P_c$ | Power used to charge the battery |
| $P_d$ | Power discharged from the battery |
| $P_{dir}$ | Power flowing directly from PV panels to meet the load or be sold to the grid |
| $P_L$ | Load |
| $P_S$ | PV panel power output |
| $S_i$ | PV generation trace for scenario $i$ |
| $\sigma$ | Standard deviation of probability distribution |
| $T_h$ | Length of time horizon |
| $T_u$ | Length of time slot |
| $\theta$ | Expected unserved energy target fraction |
| $U$ | Initial energy content of battery |

Importantly, the optimal operating strategy is typically an output of the optimization program and is dependent on the input scenario. An algorithm, or *solver*, is used to search the space of feasible solutions to find the one which maximizes (or minimizes) the objective function for the given parameters. For example, the authors of [30] formulate the problem of sizing a battery to meet the energy demands of a microgrid as a mixed-integer linear program. In [48], the problem of sizing batteries and solar panels under a fixed budget to maximize the revenue of a solar farm is formulated as a non-linear optimization problem, which is linearized to reduce the solution time.

Another notable optimization approach is to formulate a *robust optimization* problem [21], in which the objective function is optimized even when the inputs are perturbed.

Robust optimization is not included in the scope of this work. Instead, we present an alternative and simpler approach to handle uncertainties in the input parameters, described in Section 4.2.

### Simulation

System simulations are also a scenario-based sizing approach. They are versatile: a simulation program can evaluate different combinations of PV panel and battery sizes, calculating metrics such as loss of load probability (LOLP) [23], expected unserved energy (EUE), and operating cost [8]. Simulations can model the system using virtually any operating strategy, such as those proposed in [98, 79, 27, 23], and can implement complex battery models [8].

### Analytical Methods

Inspired by the analogy between energy buffering by batteries and data buffering in computer networks, a variety of analytical methods have been proposed for storage capacity sizing in the literature. For example, in [35] the system is modelled as a cyclic non-homogenous Markov chain, and the authors propose a steady-state analysis to determine whether a given system size is sufficient to meet a target LOLP. In [9], the authors use a probabilistic tail bound on the aggregate of many regulated energy demand loads to jointly size the battery capacity and transformers for a certain LOLP in a residential setting.

Among existing analytical approaches, stochastic network calculus (SNC) [68] has shown great robustness and accuracy. This approach has been used in several applications: battery sizing to reduce reliance on diesel generators in rural areas with unreliable grid connections [147], energy demand management in a fleet of electric car charging stations [100], gaining energy flexibility through heating/cooling systems in data centres [20], and supply-demand matching for prosumers [49, 171, 137].

Applying stochastic network calculus to energy systems has some subtleties, due to the unique statistical properties of the underlying energy processes and the storage model in use. This has led to a series of incremental improvements in this field of research. The idea of using stochastic network calculus for energy systems was proposed in [171], where the authors assume ideal storage devices and use affine functions to separately model the long-term behavior of each of energy demand and energy supply. In [49], the authors further develop this approach by assuming a more realistic storage model and more complicated uni-variate envelopes for energy demand and supply.

### 4.1.2 Deployment Context

Literature on sizing or analyzing the performance of PV-storage systems is extensive. Existing work may be categorized into two main areas of the context of the system's deployment: *stand-alone* or *off-grid* systems, and *grid-connected* systems. Excellent reviews of the literature can be found in [87, 88, 144]. We note that related problems are studied in the context of microgrid planning, for example in References [34, 56, 32], which look at problems such as optimizing the choice of renewable technology, and the economic feasibility of going off-grid. Such studies often feature the microgrid planning simulation model tool known as HOMER [16]. HOMER is highly customizable but does not provide a measure of confidence in the performance of the recommended system sizing, which is the focus of our work.

#### Stand-alone systems

Stand-alone systems are those in which the system can only rely on solar power and storage to meet the demand power. Several papers study the optimal sizing and cost analysis of stand-alone PV systems [82, 88, 99, 121]. The objective in designing stand-alone systems is to minimize the cost of PV and storage while meeting the power demand with an associated target LOLP or similar criteria. System cost is usually expressed either in terms of the initial capital cost [122, 145], or in terms of the annualized cost while accounting for differences in the lifetime of batteries and PV panels [55, 151].

#### Grid-connected systems

Grid-connected systems exist in different forms, for example, a residential PV-storage system, or a solar farm with supporting battery. Residential installations of PV-storage systems have the option to serve the demand by using PV panels, storage, or the grid. The price of buying or selling electricity from or to the grid can be a function of the time of the day and season, among other things. Residential installations mostly aim at selling their excess power to the grid and buying their power shortage from the grid. These options create many challenging problems that result in different objective functions.

Barra et al. [19] optimally size PV panels and storage such that a minimum target fraction of the total demand is guaranteed to be met by the PV-storage system and the cost of energy is minimized. Azzopardi and Mutale [14] minimize the annual net cost, using a case study of a residential installation where energy can be stored, used, or sold .

They consider fluctuations in time-of-use pricing and employ mixed integer programming to find the optimal size of each system component.

Using a similar model, Ru et al. [139] formulate an optimization problem to determine the critical size of the battery after which an increase in size gives no performance benefit. Other works attempt to maximize the net benefit (benefit minus cost) of a grid-connected solar PV panel with no storage [96, 97].

## 4.2 Robust System Sizing

The sizing of PV and batteries is the subject of multiple research studies. As a rule, most studies, for example [181, 72, 125, 23, 114], employ a specific and singular sizing approach based on historical solar and load data. Furthermore, from our point of view, robustness of the result to the possible and expected variations in the *future* operating environment has not been among the subjects or objectives of these studies.

In contrast, we analyze and compare multiple sizing approaches, extending them as necessary to achieve a robust result. A robust system size is one which, with sufficient probability, satisfies the system's design constraints upon reasonable[3] future fluctuations in the operating environment. In this work we consider robustness of the result as an objective.

The approaches we study use historical solar generation and electricity consumption time series as input to compute a system sizing. Given that this data is difficult to obtain for a horizon long enough to adequately capture the non-stationarity of the underlying stochastic processes, we assume that the future is statistically similar to the past and that future fluctuations in the operating environment are at least partially represented in the available historical data. This assumption is reasonable for PV generation data, where the underlying process is relatively stable (the sun), and may be reasonable for load data when there are no significant changes to the electrical appliances and their use. Nevertheless, any practical data-driven approach must take steps to prevent overfitting to historical data while making some assumptions on the future. In the context of this work, computing a robust result on the base of historical data is equivalent to resolving the problem of overfitting to this data. In other words, resolving the problem of overfitting leads us to a robust result.

[3]We believe it is unreasonable to size for *any* possible changes in load. For instance, if a homeowner buys an electric vehicle in the future, their PV-storage system size may no longer be adequate.

We extend three approaches to computing a system size: simulation, mathematical programming, and stochastic network calculus. With simulation and mathematical programming, we compute a robust size by using upper probability bounds on the sizes that meet the performance requirements on historical data. With stochastic network calculus, we essentially reduce the available data to a set of representative features that are then used to compute probability bounds on the performance targets corresponding to any given size.

The key contributions of the work presented in this chapter are:

- We provide a theoretical foundation for robust and practical sizing of both solar PV generation and storage based on simulation and mathematical programming.
- We evaluate the robustness and computational complexity of these approaches in a realistic setting.

The stochastic network calculus approach is excluded from the list of contributions because it has been developed by other members of the research team. However, we implement and evaluate this approach, as described in [74]. The program modules for computing robust PV-storage system size via simulation and stochastic network calculus have been released and are freely accessible [73].

### 4.2.1 Problem Specification

The goal of our work is to provide practical advice on how to size both solar panels and storage in a robust fashion to partly or completely offset grid usage. This section discusses the inputs and optimization criteria of the sizing problem, and how we model load as a constraint.

#### Input Parameters

It is reasonable to assume that an entity making a sizing decision would have access to a representative set of load traces, especially with the widespread deployment of smart meters that typically measure hourly load[4]. It is also possible to obtain hourly solar radiation traces in the geographical location of the entity, for most parts of the world [52],

[4]Finer-grained traces would, of course, be good to have, but unlikely to be available in practice.

and calculate the corresponding power generated from PV panels with reasonable accuracy [13] .

In keeping with prior work, we assume that these historical traces are generally representative of loads and generation. Nevertheless, the future will never exactly mimic the past; if it did, we would be able to make decisions with absolute accuracy. Thus, the sizing decision must be robust to perturbations in the inputs, i.e., to 'small' changes in the solar irradiation or loads (we elaborate on this in Section 4.2.2).

In addition to generation and load traces, we need two other parameters. Firstly, we need to know how a decision is made to either inject power into or withdraw power from the storage system. This *operating policy* can be quite complex, and is the subject of much research [98, 35, 22, 48]. Nevertheless, simple rules such as 'store excess solar energy' and 'discharge the store when solar generation is less than the load' are often adequate for many situations. We assume that, for the case of simulation and stochastic network calculus approaches, such an operating policy is provided to the sizing decision-maker. Secondly, we need a dynamic model of the behaviour of a storage system in response to power injection and discharge. For this purpose, we use the C/L/C model described in Section 3.4.

To summarize, we assume that the sizing decision-maker has access to the following information:

- A representative set of *solar traces* $S = \{S^i\}$ (for example, one trace per year; we discuss this point in greater detail in Section 4.2.2).
- A representative set of *load traces* $D = \{D^j\}$ that constitute a set of *load scenarios*. The durations of all load and solar traces should be equal.
- An *operating policy*: for the simulation and stochastic network calculus approaches, the set of rules that determine when the store is charged or discharged.
- A *storage model*, along with all associated model parameters: for a given state of charge and applied power, this is a set of equations that computes the updated state of charge.

### Criteria of optimization

Given the input parameters described in Section 4.2.1, our objective is to compute the optimal size for solar PV panels and the storage capacity. What constitutes the optimal

size will depend on the objectives of the prospective system owner. Several quality metrics or criteria are plausible[5]:

- **Minimize loss of load probability (LOLP):** This is the probability of the event that the system, where PV generation is the only source, is unable to meet the load. This probability can be numerically estimated as the ratio of the time during which the load is unmet to the total time under consideration.
- **Minimize expected unserved energy (EUE):** This is the total amount of load (energy) that cannot be satisfied by the system, based only on PV generation, during the period under consideration. If this load is not met from the grid or some other external power source, there will be user discomfort.
- **Minimize financial cost:** This is the monetary cost of purchasing the solar panel and storage system, as well as the cost of purchasing, as necessary, electricity from the grid, at its currently prevailing price. It may be viewed as a one-time capital cost added to a periodical operational expense, such as potential purchases from the grid as well as compensation for degradation of equipment from wear and tear. Note that if we could associate a cost to meeting unmet load using the grid or a local generator, then the cost-minimization objective would incorporate the objective of minimizing the unmet load.
- **Maximize robustness:** This is the degree of sensitivity of the sizing to perturbations in the input. Intuitively speaking, we wish to pick an approach such that the resulting system will meet its performance target in the future, which may have small variations from has been observed in the past. We discuss this point in greater detail in Section 4.2.2.

Computation time is also a factor in the design of the method used to compute an optimal size. We expect that the size will be chosen on behalf of a system purchaser by a sizing decision maker. The computation cost of each such decision, therefore, should not be onerous.

In many cases, there will be a trade-off between cost on the one hand, and LOLP/EUE and robustness on the other. Moreover, robustness and computation cost are closely associated, since getting robust results generally implies the processing of greater amounts of data. In this work, we focus on minimizing the cost of solar PV and storage, subject to meeting a certain LOLP or EUE constraint.

[5]For each application, one or multiple of these items can serve as objectives and one or multiple others as constraints.

#### System Performance Metrics

Traditionally, the LOLP/EUE target is specified together with a length of time over which this criterion should be met [168]. For example, a common loss-of-load target for reliable grid-scale electrical systems is one day over a period of 10 years, corresponding to an LOLP target of 0.000274. Such a high level of reliability makes sense where frequent or prolonged loss-of-load events correspond to millions of dollars in losses to the economy supported by the electrical system. Achieving this level of reliability equates to sizing for the worst-case behaviour with virtually 100% confidence, requiring expensive systems that are oversized for the average behaviour but are nevertheless cheaper than the cost of loss-of-load events.

In contrast, smaller systems such as a house may typically tolerate higher LOLP due to smaller penalties associated with loss-of-load events, the ability to easily shut off electrical appliances at times of high load, and the availability of the grid and perhaps a local generator to offset some of these events. For these reasons, an excessively pessimistic sizing of the system with an LOLP of $\approx 0$ over long periods with 100% confidence would be a sub-optimal financial decision. In these conditions, it would be appropriate for a home-owner to choose a more affordable target, for example an LOLP of 5% over all 90-day periods, with a confidence of 95%. Mathematically, this corresponds to a system which guarantees $\mathbb{P}(LOLP \leq 0.05) \geq 0.95$ over any 90-day period. In this work, we refer to the combination of time interval, LOLP/EUE target, and confidence, as the quality of service (QoS) target.

### 4.2.2 The Impact of Non-Stationarity

One of the key insights in this work is that the traces which serve as input to any sizing approach may neither be stationary nor definitively representing the future. We discuss this next.

#### Traces, Trajectories, and Stochastic Processes

A solar or load trace with $T$ entries of the form (*time*, *value*) is a trajectory instantiated from a stochastic process, which is defined as a set of random variables indexed by time. That is, $S^i(t)$, the $t^{th}$ element of the $i^{th}$ solar trace (resp. $D^j(t)$, the $t^{th}$ element of the $j^{th}$ load trace) is a value assumed by the random variable $S(t)$ (resp. $D(t)$) from a corresponding distribution. Hence, we can fully characterize the *historical* solar (resp. load) stochastic process by defining a joint distribution of a set of $T$ random variables,

one for each time step. If we assumed independence of each time step, we could decouple these distributions. This would allow us to use the set $S$ (resp. $D$) of solar generation (resp. load) traces to estimate parameters for each of the $T$ distributions. For example, the numerical mean of the $t^{th}$ time step of the set of traces can be viewed as an estimate of the mean of the $t^{th}$ distribution and the sample variance of this set is an estimate of its variance. Thus, with sufficient data, we can use standard statistical techniques to find the best-fitting distributions that characterize a set of traces.

Given this characterization of historical stochastic processes, what can we say about their future states? Suppose that the generation and load stochastic processes are time-invariant. Then, once the historical processes are characterized, the future is also 'known' in the sense that we may generate potential future trajectories by generating a random value per time step from the corresponding distribution. We can then choose a sizing that meets our sizing objectives not just for historical trajectories, but also for potential future trajectories.

However, this naive approach has three problems. Firstly, even assuming independence of time steps, it is onerous to define $T$ separate distributions, since $T$ can be very large, on the order of 10,000 – 100,000 values. Secondly, there is no guarantee that a stochastic process parametrized with historical traces will adequately represent the future. Thirdly, we do not have any definite criteria on how much data would be sufficient for parameterizing these processes. To resolve these problems, a closer look at the generation and load stochastic processes is required.

### Causes of Non-Stationarity

A key observation is that both the solar and load stochastic processes are *non-stationary*[6] due to the following reasons:

1. **Diurnality.** For example, the distribution of the r.v. $S(t)$ corresponding to a time slot $t$ at night will differ from the distribution of an r.v. corresponding to a time slot at mid-day.

2. **Seasonality.** For example, the distribution of the r.v. $S(t)$ corresponding to a time slot $t$ at mid-day in winter will differ from the distribution of an r.v. corresponding to a time slot at mid-day in summer.

[6]Generally speaking, this means that statistics computed from two different random sub-samples of the traces can differ.

3. **Long-term trends.** For example, the distributions of the r.v.'s $D(t)$ and $S(t)$ corresponding to a time slot $t$ at the start of a trace may differ from their distributions for a time slot later in the trace.

4. **Autocorrelation.** For example, the distributions of the r.v.'s $S(t)$ and $D(t)$ are dependent on the values taken by the respective r.v.'s $S(t-1)$ and $D(t-1)$.

Non-stationarity should be taken into account upon characterizing historical generation and load stochastic processes.

### Stochastic Process Parametrization

Recall that the parameters of the stochastic process, i.e., corresponding to each of the $T$ distributions constituting the process, are derived from solar and load traces. Given that the process has both diurnal and seasonal non-stationarity effects, the solar and load traces must be both *detailed* enough and *long* enough to capture the effects. More precisely:

- The traces should have sufficient temporal resolution to capture diurnal changes. That is, the time step should be sufficiently small in order to have an adequate number of values for each part of the day.
- The traces should be long enough to capture seasonality, i.e., at least one year in duration, if not longer.
- The traces should be long enough to capture any long-term trends in load. We assume that PV generation from one year is statistically similar to PV generation from another year, for the same location.
- There should be enough traces in the set of traces so that there would be sufficient samples for adequate estimation of the parameters of each distribution.

Ideally, we would have access to per-minute or even per-second load and generation traces spanning several decades. Then, setting $T = 60 \times 24 \times 365 = 525,600$, i.e., per minute of the year, we would obtain multiple sample values for each time step, allowing us to estimate, with adequate confidence, the parameters of each of the $T$ solar generation distributions, and potentially long-term trends in the load (for example by fitting a linear regression to the residual after accounting for diurnal and seasonal effects).

In practice, it is unlikely that such data traces are available. Load and solar generation are often measured with a frequency of once every 30 minutes, if not lower frequency, and

it is extremely rare that more than a few years of data are available. Therefore, we resort to the following pragmatic approach that, in our experience, works reasonably well. Given a dataset in the form of a time series that is at least one year long, and begins and ends on the same day of the year, we join the ends of the data to create a circular time series. We then sample blocks of $X$ days, where $X$ is the time interval in the QoS target, from the time series by choosing random start times. We call each of the $n$ samples a 'scenario' and treat their ensemble as an estimate of expected future scenarios.

We now discuss three robust sizing approaches that base their sizing decisions on scenarios created in this way. We defer a discussion on how to evaluate the robustness of these approaches to Section 4.2.4.

### 4.2.3 Robust Sizing Approaches

The sizing decision, that is, choosing the size of the store $B$ (in kW or kWh) and of the solar panels $C$ (in kW) to meet one or more of the objectives discussed in section 4.2.1, can be made using many different approaches. In this section, we explain three representative approaches.

We make the following assumptions:

- For simplicity, we assume that the goal is to find the minimum-cost storage and solar PV sizes that meet a certain LOLP or EUE QoS criterion.
- We assume that we have available a solar generation trace $S$ and a load trace $D$ corresponding to the same time interval with a length of at least one year. As already discussed, from these traces we can obtain $n$ scenario samples of $X$ days of solar generation and the corresponding load.
- We size the storage system for only energy, not for power, since sizing for power is typically trivial (the power rating of the storage system must exceed the sum of power draws of the set of simultaneously active load components).
- We assume the storage system energy capacity can only take one of $b$ different values and that the solar panel size can only take one of $c$ different values. Subsequently, sizing a system is essentially conducting a grid search through $b \times c$ pairs of solar PV panel sizes and storage capacity sizes in order to determine the optimal size, i.e., the size with the minimum cost which guarantees the satisfaction of the QoS.

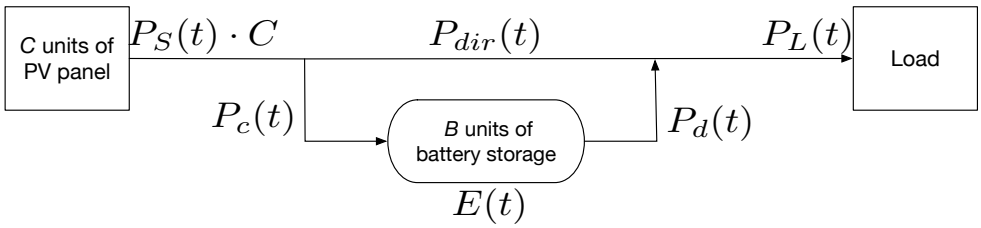


Figure 4.1: Off-grid PV-storage system diagram

- We assume that if a certain combination of storage and PV values results in a certain LOLP/EUE, then larger values of either storage or PV will always result in lower values of LOLP/EUE. This allows us to use a greedy grid-search heuristic.
- We assume that the operating policy is given by Algorithm 1, which maximizes consumption of PV generation and is described in Section 5.2.1. We refer to it as the *simple* operating policy.

We denote the number of time steps in the load and solar generation traces by $T$. $\pi_B$ is the price for one unit of battery (i.e., 1 cell, or one module), and $\pi_C$ is the price for one unit of PV panel. We normalize the solar generation $S^i$ trace, so that it represents the power generation from a single unit of PV panel, which we denote as $P_S$. A demands trace $D^i$ represents the household electricity load, which we denote as $P_L$. Finally, the LOLP target is denoted $\epsilon$, and the EUE target, expressed as a fraction of the total load, is denoted $\theta$.

### Optimization

In this approach, we formulate an optimization program for solar panel and battery sizing with the objective of minimizing the capital cost of the system, subject to physical system and LOLP/EUE constraints. We do not specify the operating policy, leaving this decision to the optimization solver, because even the simple policy (Algorithm 1) is difficult to encode in a mathematical program. The optimal policy may not be practical since it does necessarily not correspond to a known set of operating rules. However, this allows us to compute the best possible size in the case of optimal operation. In this sense, although the sizing decision made by the optimization program is a potentially-unattainable lower bound, it measures the level of sub-optimality in the operating policy used in the two other approaches.

Our approach has two phases. In the first phase, for each scenario and for each of the $b$ potential battery sizes, we compute the optimal solar panel size $C$, assuming optimal

operation. This gives us $n$ *sizing curves* defined by the interpolation of $b$ sizes computed for each scenario (see Figure 4.2). In the second phase, we use a technique based on the Sample Univariate Chebyshev bound [141] to compute a *robust* size that is insensitive to the details of individual traces. We discuss each phase in turn.

***Phase 1***

For phase 1, define $P_c$ to be the charging power, $P_d$ to be the discharging power, $P_{dir}$ to be the power that flows directly from PV panel to load, and $E$ to be the energy content. The number of battery cells is $B$ and the number of PV panel units is $C$. Figure 4.1 shows a labelled system diagram.

The battery model used here is the C/L/C from Section 3.4 with the following parameters: $\eta_c$ (resp. $\eta_d$) the charging (resp. discharging) efficiency, $\alpha_c$ (resp. $\alpha_d$) the charging (resp. discharging) rate limit, $u_1$, $v_1$, $u_2$, $v_2$ used to characterize the power-dependent lower and upper limits on the energy content (see constraint (4.6)). The model parameter values are calibrated for a single cell, and the charging/ discharging rate and energy content limits are scaled with the number of cells, $B$. The energy content at the end of time slot $t$ is denoted $E(t)$, and the initial energy content is $U$. The duration of a time-step is $T_u$ and the number of time-steps in a data trace is $T$.

We first present the full formulation with an LOLP constraint, and then show how to modify it for an EUE constraint. Given a scenario $(P_S(t))$, $(P_L(t))$, and storage parameters $B$, $\eta_c$, $\eta_d$, $\alpha_c$, $\alpha_d$, $u_1$, $v_1$, $u_2$, $v_2$, $U$, and trace parameters $T_u$ and $T$, the problem can be formulated as:

$$\min_{\substack{C,P_c,P_d,\\ P_{dir},I,\gamma,E}} C \tag{4.1}$$

subject to

$$P_c(t) + P_{dir}(t) \leq P_S(t)C \quad \forall t \tag{4.2}$$
$$P_{dir}(t) + P_d(t) = P_L(t) - \delta(t) \tag{4.3}$$
$$E(0) = U \tag{4.4}$$
$$E(t) = E(t-1) + P_c(t)\eta_c T_u - P_d(t)\eta_d T_u \quad \forall t \tag{4.5}$$
$$u_1 P_d(t) + v_1 B \leq E(t) \leq u_2 P_c(t) + v_2 B \quad \forall t \tag{4.6}$$
$$0 \leq P_c(t) \leq B\alpha_c \quad \forall t \tag{4.7}$$
$$0 \leq P_d(t) \leq B\alpha_d \quad \forall t \tag{4.8}$$
$$I(t) \in \{0, 1\} \quad \forall t \tag{4.9}$$
$$C, P_{dir}(t), \delta(t), E(t) \geq 0 \quad \forall t \tag{4.10}$$
$$1/T \sum_{t=1}^{T} I(t) \leq \epsilon \tag{4.11}$$
$$I(t) \leq \delta(t)Z \quad \forall t \tag{4.12}$$
$$\delta(t) \leq I(t)P_L(t) \quad \forall t \tag{4.13}$$
$$P_c(t)P_d(t) = 0 \ \forall t \tag{4.14}$$

Constraint (4.2) states that the sum of what goes into the battery and directly towards the load is bounded by the solar generation. $\delta(t)$ is the load that is not met from solar generation at time $t$ (it is always $\leq P_L(t)$), constraints (4.5)–(4.8)) represent the battery model. $I(t)$ is an binary variable used to indicate if the load is met or not in time-step $t$ ($I(t) = 1$ means the load is not met). Constraint (4.12) ensures that $I(t)$ is zero if $\delta(t) = 0$ ($Z$ is a large positive constant), and Constraint (4.13) ensures that $I(t)$ is one if $\delta(t) > 0$. Constraint (4.11) is the LOLP constraint. Constraint (4.14) forbids simultaneous charging and discharging, though it was shown in [48] that it can be ignored which makes the problem an Integer Linear Program (ILP). Note that in this program $B$ and $C$ are real numbers, i.e., they are not limited to the pre-defined values used for the other two approaches.

To express an EUE constraint, this problem formulation can be modified as follows.

Replace Constraints (4.11)–(4.13) with:

$$\sum_{t=1}^{T} \delta(t) \leq \theta \sum_{t=1}^{T} P_L(t) \tag{4.15}$$

Note that the formulation with the EUE constraint is a linear program (LP), which can be solved much more efficiently than an ILP.

We use this mathematical program to compute the smallest[7] $C$ for each of the $b$ battery sizes so that the system meets the QoS target; these points define a curve in a $(B, C)$ space. We denote each curve as $K_i$ corresponding to the $i$th scenario, for a total of $n$ curves.

***Phase 2***

In phase 2, we use the $n$ sizing curves obtained in phase 1 to compute a probability bound on the system size with a given measure of confidence.

First, for each of the $b$ values of $B'$, we construct the set $L_{B'}$ consisting of points in the $(B, C)$ space along the intersection of the line at $B = B'$ and each curve $K_i$.

$$L_{B'} = \{C'' : C'' = K_i(B')\} \tag{4.16}$$

Each set of points can be viewed as samples from a distribution defined by the sizing curves. Denote the size of the set $|L_{B'}| = N_{B'}$. Not all sizing curves have a value defined at $B'$, so $N_{B'} \leq n$. We can compute a sample Chebyshev bound, as formulated in [153], on the $C$ values as follows:

$$\begin{aligned} &\mathbb{P}\big\{|C - \mu_{C,N_{B'}}| \geq \lambda \sigma_{C,N_{B'}}\big\} \leq \\ &\quad \min\Big(1, \frac{1}{N_{B'}+1}\Big\lfloor \frac{(N_{B'}+1)(N_{B'}^2 - 1 + N_{B'}\lambda^2)}{N_{B'}^2 \lambda^2} \Big\rfloor\Big) \end{aligned} \tag{4.17}$$

The inequality above is a bound on the probability that the distance between some future value of $C$ for the corresponding $B'$ from the estimated mean $\mu_{C,N_{B'}}$ exceeds a factor $\lambda$ of the estimated standard deviation $\sigma_{C,N_{B'}}$. To use this inequality to compute a size, we first find the smallest $\lambda$ that satisfies our confidence measure $\gamma$:

$$\min_{\lambda}\Big(\frac{1}{N_{B'}+1}\Big\lfloor \frac{(N_{B'}+1)(N_{B'}^2 - 1 + N_{B'}\lambda^2)}{N_{B'}^2 \lambda^2} \Big\rfloor\Big) \leq 1 - \gamma \tag{4.18}$$

[7]The solution gives us the optimal $C$ as a real number, which we round up to the nearest potential $C$ value among the $c$ possibilities.

Next, we rearrange the inequality in the LHS of Eq. (4.17) to obtain a robust value $C^*_{B'}$ using the $\lambda$ that satisfies Eq. (4.18):

$$C^*_{B'} = \mu_{C,N_{B'}} + \lambda\sigma_{C,N_{B'}} \tag{4.19}$$

The resulting set of points $(B', C^*_{B'})$ can be interpolated to define a curve which we call the *Chebyshev curve on C*, since each point on the curve is a Chebyshev bound on $C$ values. Similarly, we can construct a *Chebyshev curve on B* by computing Chebyshev bounds on the following sets for each of the $c$ values of C':

$$L_{C'}\{B'' : Z_i(B'') = C'\} \tag{4.20}$$

The upper envelope of these Chebyshev curves represents system sizes which are *robust* with respect to both $B$ and $C$ with confidence measure $\gamma$. We use the least-cost system along the upper envelope as the final recommendation for the robust size.

If we are confident that the estimated mean and standard deviation have converged to the population mean after $n$ samples, we can obtain a tighter bound characterized by a value of $\lambda$ that satisfies the following:

$$(1 - \lambda)^{-1} = 1 - \gamma \tag{4.21}$$

Using Eq. (4.21) in place of Eq. (4.18) corresponds to the classical *analytical* Chebyshev bound which assumes that the population mean and standard deviation are known.

***Computation cost***

The two Chebyshev curves can be computed with $O(bn + cn)$ computations, hence the computation time of this approach is dominated by the computation of the sizing curves via the optimization program. The inputs to the optimization program are the solar and load traces, each of size $O(T)$, for a total size of $O(T)$. Asymptotically, this is also the number of variables in the program. Denoting by $Q = O(T)$ the number of variables and $L = O(T)$ as the number of bits of input to the algorithm, even for an LP, which is far more computationally efficient than an ILP, the best-known approach, the interior-point method, requires a runtime of $O(Q^{3.5}L^2 \cdot logL \cdot loglogL) \approx O(T^{5.5} \cdot logT)$ [167]. Since we need $nb$ such runs, and our problem is integer, the total complexity is lower bounded by $O(nbT^{5.5} \cdot logT)$.

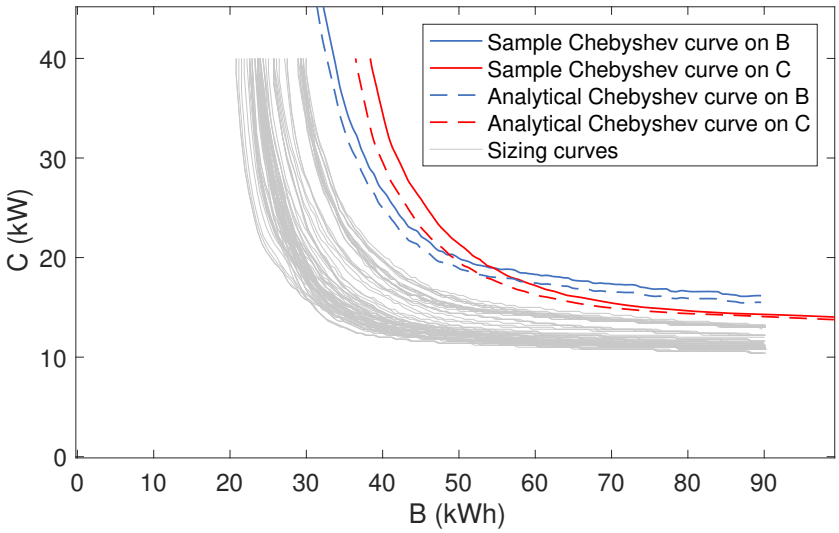


Figure 4.2: Fifty sizing curves and corresponding sample and analytical Chebyshev curves for $\gamma = 0.95$.

### Simulation

In this approach, we run system simulations to construct the sizing curves for each of the $n$ scenarios. Specifically, for each scenario, for each potential sizing choice, and for each time step $t \in [1, T]$, we determine the availability of solar power $S(t)$ and the load $D(t)$. Depending on these values, the storage and PV sizes under test, and the given operating policy, we use the storage model to either charge or discharge the store, updating its energy content as in Eq. (4.5). If we find that there is a need to discharge the store, but its SoC is zero, then we mark that the load is unmet from solar generation for this time step. At the end of each simulation, we empirically determine the LOLP $\epsilon$ or EUE $\theta$ for this sizing.

We use a search algorithm that, given a scenario, efficiently searches the $(B, C)$ space to compute the sizing curve:

**Step 1**: For $C = C^{max}$, begin at $B = B^{min}$ and increment $B$ to find the smallest value of $B$ such that the system satisfies the target performance requirement; record $(B, C)$.

**Step 2**: Decrement C, then start from the most recent value of $B$ and increment it until the system satisfies the performance requirement; record $(B, C)$.

**Step 3**: Repeat previous step through $C = C^{min}$. Recorded $(B, C)$ pairs compose the sizing curve.

This algorithm first finds the edge of the curve at $C = C^{max}$, then traces the rest of the curve using at most $b + c$ simulations. We then use phase 2 as described in Section 4.2.3 to process these curves and compute a robust size from the upper envelope of the two Chebyshev curves.

Note that the computation cost of this approach is $O(nT(b + c))$, since each time step takes $O(1)$ computation time, there are $T$ steps per simulation, and $(b + c)n$ simulations.

#### Stochastic Network Calculus (SNC)

Unlike the mathematical programming and simulation approaches, which can compute either LOLP or EUE metrics with minor changes, stochastic network calculus has significant differences in the mathematical formulation used to compute the sizing for each metric. Given a sizing, a QoS target and a set of scenarios, the SNC approach computes an upper-bound on the LOLP or EUE in each scenario. We consider a sizing to be valid if the percentage of scenarios that meet the LOLP or EUE target is within the confidence target.

The formulation of this approach is not provided in this work; it was developed in [74] by the team of researchers where the author of this thesis is responsible only for the software implementation. The results of the implementation are presented in the following section. The computational complexity of the approach is $O(nT(b + c))$.

### 4.2.4 Numerical Evaluation

A numerical evaluation of our approaches is based on four years of PV generation and load data collected from homes in the Pecan Street Dataport [1]. A detailed view of results for three homes with low, mid, and high levels of consumption, as well as an aggregated sizing evaluation across 52 homes from this dataset, are presented.

To evaluate the cost of a particular size, we set the installed cost of solar panels $\pi_C$ to be USD 2.50/W, and the cost of storage $\pi_B$ to be USD 460/kWh[8], with battery parameters corresponding to a Lithium-Nickel-Manganese-Cobalt (NMC) battery chemistry [92, 38] as summarized in Table 4.2. We use the simple operating policy of charging the battery when solar generation exceeds the load and discharging the battery when load exceeds solar generation. At the beginning of the process, the battery is in a fully charged state.

Although our optimization-based sizing approach is robust, in that it is insensitive to small perturbations in the input trace, it is often impractical to use this approach; it relies

[8]Source: https://www.tesla.com/powerwall

Table 4.2: NMC battery model parameters

| **Parameter** | $\alpha_c$ | $\alpha_d$ | $u_1$ | $u_2$ | $v_1$ | $v_2$ | $\eta_c$ | $\eta_d$ |
|---|---|---|---|---|---|---|---|---|
| **Value** | 1 | 1 | 0.053 | -0.125 | 0 | 1 | 0.99 | 1.11* |

*includes inverter inefficiencies of ∼10%

on optimal operating policy which cannot be determined in advance. Thus, we evaluate our Chebyshev-curve sizing approach using sizing curves computed via simulation only, and compare the resulting size with the one computed using the SNC approach.

### Convergence of result

In accordance with the simulation and optimization approaches, the recommended system size is based on a statistical measure of the underlying samples of computed sizing curves. With SNC, we compute a probabilistic upper-bound on the number of loss-of-load events for each scenario in the ensemble. With all three approaches, the recommended size is progressively refined as more scenarios are evaluated.

We find that in all three approaches, mean $B$ and $C$ values converge after about 100 randomly sampled scenarios for $\epsilon$ or $\theta$ targets of 0.05. For smaller targets, more scenarios are required for convergence.

### Sizing

Figure 4.3 shows the robust system sizes computed using SNC and simulation approaches for $\epsilon = 0.05$ and $\theta = 0.05$, with a confidence of 95% and 97%. The figures include the sizing curves as well as the Chebyshev curves computed from them. The sizing obtained for all approaches with 95% confidence lie above the sizing curves. Greater values of the confidence parameter lead to a more pessimistic sizing, which is more robust to variations in scenarios that might be observed in the future but not fully reflected in the historical data.

We note that the size calculated using the optimization approach (not included in Figure 4.3) is identical to the size obtained using simulations for an EUE target. However, for an LOLP target, the size computed using the optimization approach is always smaller than the size computed with the other two approaches. This is expected because the optimization approach uses an optimal operating policy, whereas the other two approaches use the simple policy (Algorithm 1). The simple policy happens to be optimal for minimizing EUE, but not for minimizing LOLP.

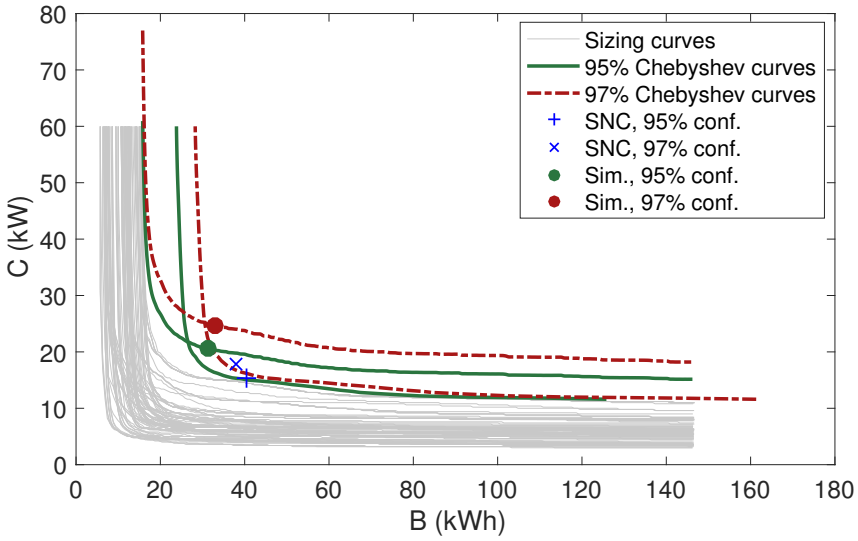


(a)

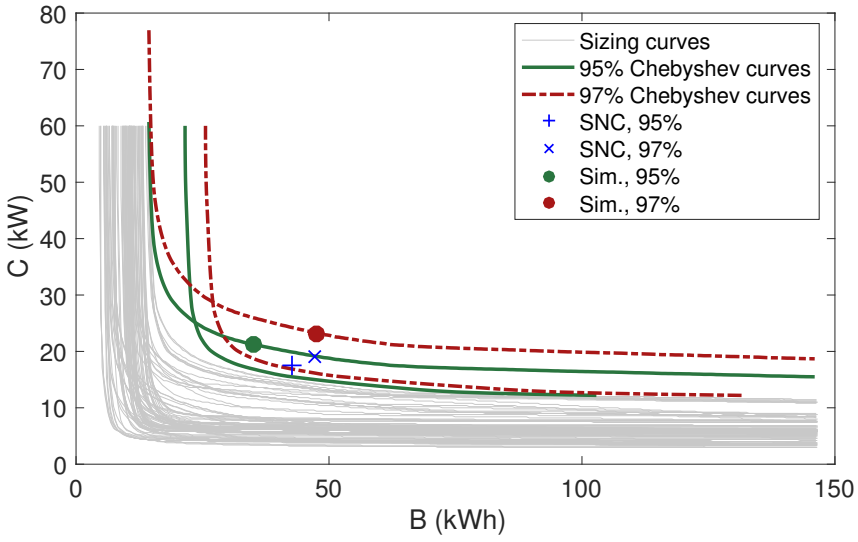


(b)

Figure 4.3: Comparing sizings at different confidences with SNC and simulation approaches, for LOLP with $\epsilon = 0.05$ (4.3a) and EUE with $\theta = 0.05$ (4.3b).

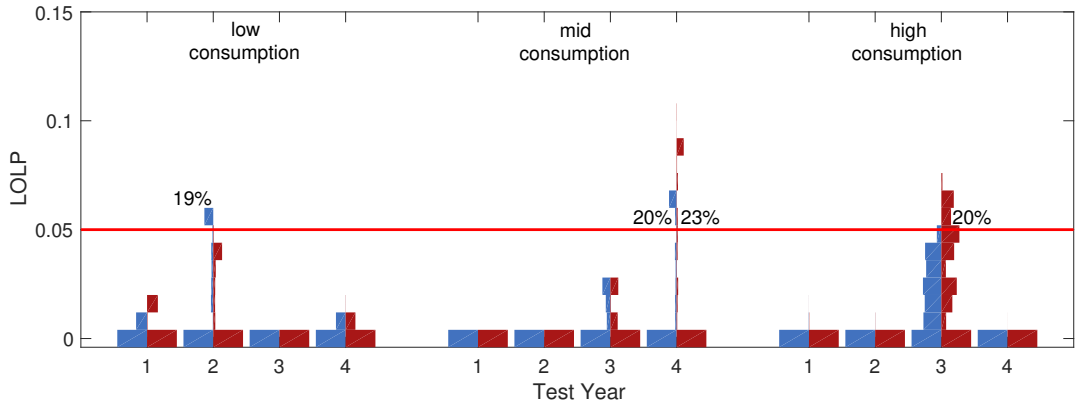


Figure 4.4: LOLP corresponding to robust system size computed for $\epsilon = 0.05$ and $\gamma = 0.95$. Left (blue) histogram for each test year is for simulation results, right (red) histogram is for SNC.

### Robustness

We compare the robustness of the size computed using the simulation and SNC approaches in Figures 4.4, 4.5, 4.6 and 4.7, which summarize the results of leave-one-year-out analysis. For each of the four years in the dataset, we randomly sample 100-day scenarios from the other three years. We then use the simulation and SNC approaches to compute a robust system size using either 95% or 97% confidence bounds for $\epsilon = 0.05$ (Figures 4.4,4.5) or $\theta = 0.05$ (Figures 4.6,4.7). Each size is tested on scenarios from the test year, and the distribution of resulting LOLP and EUE values for each scenario is presented as a histogram. The vertical histograms show the leave-one-year-out sizing performance on scenarios sampled from the test year for three households, four years for each. The resulting LOLP or EUE of the system with size computed using the simulation-based approach is shown in the blue histograms extending to the left, while red histograms extending to the right show comparable results with the SNC approach. Figures 4.4 and 4.5 are for an LOLP target of 0.05, with 95% and 97% confidence respectively. Figures 4.6 and 4.7 are for an EUE target of 0.05, with 95% and 97% confidence respectively. For test years having scenarios where the LOLP/EUE exceeds the target, there is an annotation showing the percentage of scenarios that exceed the target.

Note also that each of the four subsets of three years of data can result in substantially different sizes. If a year with particularly high load is left out, such as the third year in the high-consumption household, the computed size results in a violation of the QoS,

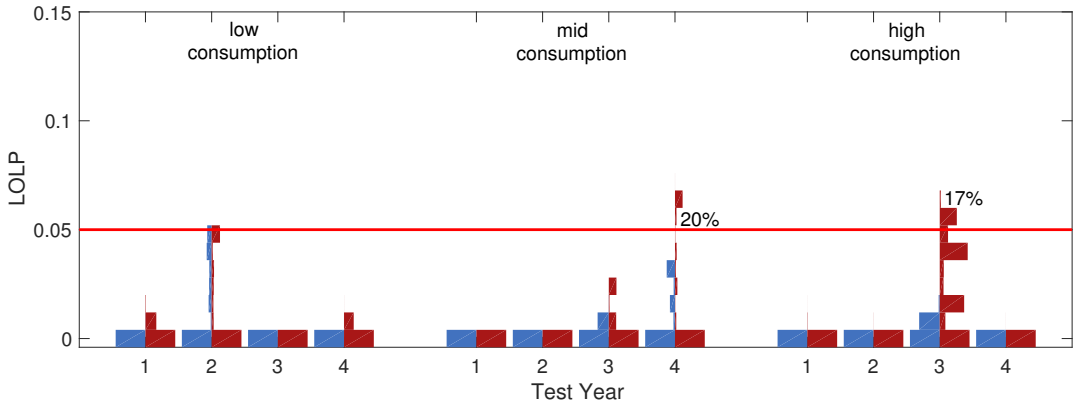


Figure 4.5: LOLP corresponding to robust system size computed for $\epsilon = 0.05$ and $\gamma = 0.97$. Left (blue) histogram for each test year is for simulation results, right (red) histogram is for SNC.

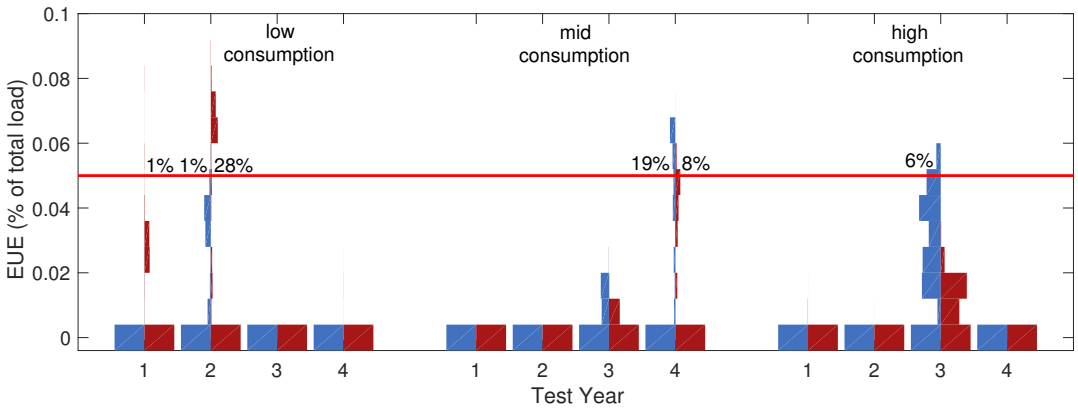


Figure 4.6: EUE corresponding to robust system size computed for $\theta = 0.05$ and $\gamma = 0.95$. Left (blue) histogram for each test year is for simulation results, right (red) histogram is for SNC.

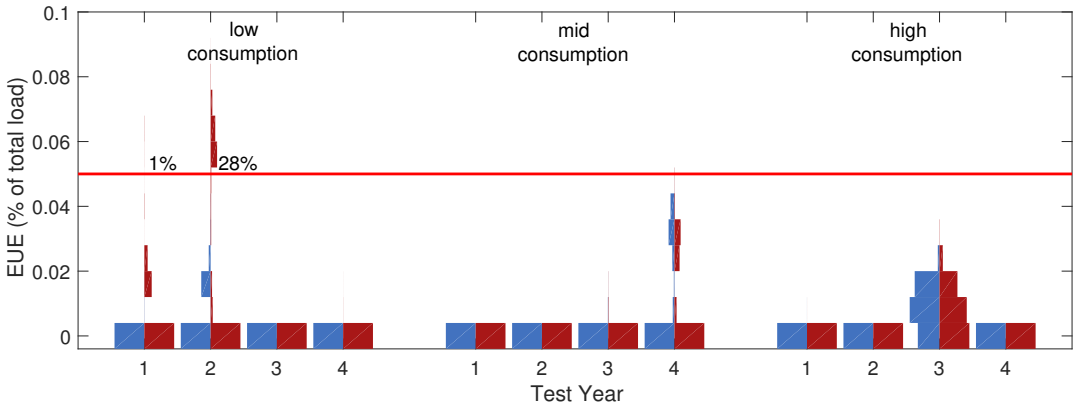


Figure 4.7: EUE corresponding to robust system size computed for $\theta = 0.05$ and $\gamma = 0.97$. Left (blue) histogram for each test year is for simulation results, right (red) histogram is for SNC.

indicating that solar generation and load may highly vary across years. Their variability may be accounted for by using a higher value of the confidence parameter. Specifically, note that with $\gamma = 0.95$, there are several instances where the sizes computed using both simulation and SNC approaches fail to meet the performance bound. This is because of atypical behaviour in one of the years compared to the other years. For both approaches, when $\gamma$ is increased to 0.97, the number of violations decreased. Compared to the SNC approach, the simulation approach is more sensitive to increases in $\gamma$, since the Chebyshev curves give very loose bounds at high confidence, while the SNC approach uses an empirical confidence measure. In practice, we expect $\gamma$ to be a user-supplied parameter that reflects their level of optimism.

The aggregated results for 52 homes from Austin, Texas with data through years 2014-2017 are shown in Figures 4.8 and 4.9. For each house, a leave-one-year-out sizing for an LOLP or EUE target of 0.05 over a period of 100 days with 95% confidence is computed. This size is then tested on 200 randomly selected 100-day periods from the test year, for a total of $52 \times 4 \times 200 = 41600$ tests for each sizing approach. The results are presented in histogram form. Notably, the fraction of values that are within the 5% LOLP and EUE target are well within $\gamma = 0.95$.

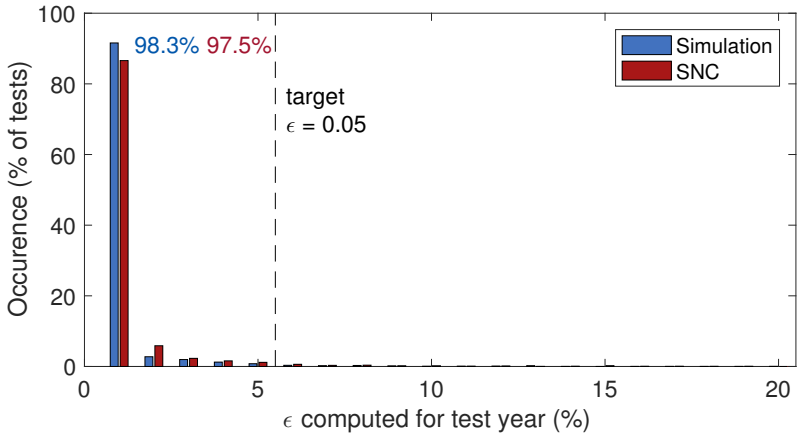


Figure 4.8: Aggregated leave-one-year-out test results on 52 houses, with LOLP target of 0.05 and confidence target of 95%. Percentage of tests that land within the LOLP target for simulation (left) and SNC (right) sizing approaches is annotated.

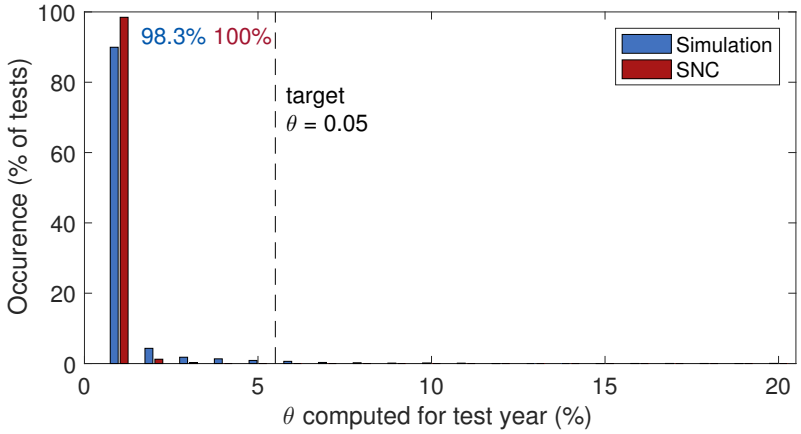


Figure 4.9: Aggregated leave-one-year-out test results on 52 houses, with EUE target of 0.05 and confidence target of 0.95. Percentage of tests that land within the EUE target for simulation (left) and SNC (right) sizing approaches is annotated.

Table 4.3: Computation time (Linux user time) for robust sizing methods

| Method | Mean CPU time for 100 scenarios (h:m:s) | | | |
|---|---|---|---|---|
| | LOLP | | EUE | |
| | $\mu$ | std. error | $\mu$ | std. error |
| Simulation | 0:0:38 | $<$ 0:0:01 | 0:0:38 | $<$ 0:0:01 |
| Optimization | 46453:20:05 | 896:6:55 | 277:45:21 | 1:01:22 |
| SNC | 0:0:24 | $<$ 0:0:01 | 0:0:16 | $<$ 0:01 |

#### Computation time

As discussed in Section 4.2.3, the asymptotic complexity of the optimization approach is lower bounded by $O(nbT^{5.5}logT)$. The computational complexity of simulation and SNC is $O(nT(b+c))$. Thus, for large values of $T$, which are typical, the best approaches are simulation and SNC, with SNC up to a factor of 2.5 faster than simulation.

Table 4.3 shows the CPU time required to process 100 scenarios, each of which has 365 days of data, on a 2.7 GHz Intel Xeon CPU. Compared to simulation and SNC approaches, optimization takes five to seven orders of magnitude more time[9] to compute a result.

### 4.2.5 Discussion

#### Comparison of the three approaches

Our work evaluates three distinct approaches to robust and practical sizing of solar and storage systems. Over and above the numerical comparison in Section 4.2.4, we now make some qualitative observations about their relative merits.

Unlike some prior works [41, 71, 48] which study the joint problem of optimal sizing and optimal operation, in this work, we study only sizing. However, with optimization, the operation rules are usually left as a free variable, in that the output of the optimization program is also the optimal charge/discharge schedule. However, these operation rules may not always be used in practice, because the rules depend in detail on the traces, and the details of the future are unknown. If we could encode operating rules into the optimization program, we would be able to calculate a size that does not have this dependence between the rules and the input data. Unfortunately, it is non-trivial to encode arbitrary operating rules in an optimization program. For instance, consider the rule "charge the store from

[9]We used CPLEX 12.6.3 solver, which has highly parallelized LP and ILP solvers

the grid if the battery charge is below 30% and the grid price is lower than \$0.12/kWh." This rule defines a dependency between the charging power and the battery charge, which complicates the formulation of the program and also makes it non-linear and hence more difficult to solve efficiently.

There is a similar problem with stochastic network calculus, where encoding complex charge/discharge operating rules may result in substantial complication of the subsequent analysis. In contrast, the simulation approach may be used with any operating strategy. Moreover, it has acceptable computation time (though greater than stochastic network calculus). Thus, from a qualitative perspective, the simulation approach is perhaps the best one, especially when combined with a Chebyshev bound.

#### Summary on robust system sizing

In conclusion to this study of robust system sizing, we highlight several contributions. Firstly, it is the first work, to our knowledge, that provides robust and practical advice on sizing through a comparison of multiple approaches. Secondly, the application of a univariate Chebyshev curve in the calculation of a robust PV-storage system size is innovative, and may be generalized to other robust optimization problems.

## 4.3 Future Work: Sizing with Limited Data

As many similar approaches, the system design approach described in Section 4.2 depends on the availability of data. This data consists of historical hourly measurements, and is used as a proxy for the future operating environment. In many cases, the availability of such data is very limited. Smart electricity meters that are capable of data collection are not yet ubiquitous, and, even if they were widely available, it would take years to collect a sufficient quantity of data.

In our work, we have assumed that the quantity of data may be sufficient for extracting an accurate statistical measure (Chebyshev bound) which is used to determine an effective system design. We have not considered the more limiting but real scenarios where the only available load data is in the form of a monthly aggregate measurement from an electricity bill, or the more extreme case where the building is new and there are no specific measurements available. In this section, we discuss possible approaches to system design in scenarios with these limitations.

According to the trends in the deployment of smart meters and other measurement devices, data limitations will not be a problem in the distant future. However, currently, the lack of data is a real limitation. Waiting until the problem is resolved naturally by the proliferation of smart meters is not consistent with the urgency in the responsible attitude to the alarming climate changes. Hence, let it be even a temporary solution, solving the practical issues in the effective design of these systems is of importance.

### 4.3.1 Aggregate data

Aggregated data represents the reduction of some range of variation into a single value. Even hourly measurements are already an aggregation of processes that exhibit dynamics at a sub-hourly level, though in the context of PV-storage systems it has been shown in [95] that hourly measurements are sufficient for system sizing under loss-of-load constraints. In the analysis of PV-storage systems, where the battery's primary functions include acting as a buffer between variations in PV generation and electricity consumption, relying on data aggregated beyond one hour prohibits us from adequately accounting for the dynamics of PV and load processes.

A possible solution to this problem is explored in [156], where the authors assume the availability of a year's worth of monthly electricity consumption data, and a database of hourly measurements from different but similar electricity consumers. The monthly data is used to find a match in the database with similar monthly consumption patterns, which is then used in the parameterization of a data generation model. The parameterized model is used to generate data which is then directed as input for the simulation-based robust sizing approach discussed in Section 4.2. This method was numerically evaluated and validated on a set of houses from the Pecan Street Dataport [1].

### 4.3.2 No data

System sizing without any supporting PV generation and load data has not been a subject of this work. Accordingly, we do not present any complete solutions to this problem. The following considerations, however, may be used as a starting point in attempts to resolve this problem.

For a user who has decided to integrate a PV panel and a battery into their local energy system, the first question that may come to mind is: what should be the size of the system? Obviously, the optimal size is always preferable, especially for a user making a large investment. A likely scenario is that the user does not have the necessary data to

determine the optimal system size using conventional data-driven optimization techniques [30, 157, 174], either because the data has not been collected, or because the system is new and has no data.

The user might have information to group themselves into a class of similar users, using data such as estimated annual electricity consumption, socio-economic status, and projected PV generation in their locale. There may be sufficient data for the class of users to calculate the system size that is close to the optimal for the class, which is a rule-of-thumb approach that gives the user a starting point for determining a suitable system size. There are several problems with this approach. Firstly, once the system is deployed, it is unlikely that its initial size will turn out to be the optimal one. Secondly, the system operating environment may change over time, for example, in electricity consumption patterns or grid pricing scheme; the optimal system size is known to be sensitive to such variations. These problems motivate a *dynamic system size* that can be adjusted to meet the changing needs of the user.

The concept of a dynamic system size is based on the idea that the optimal system size can be computed after the system is installed and data on the operating environment is collected. To bring this idea to reality, changing the capacity of the PV panels or battery would need to be an easy task. Although both components are modular, which presents an opportunity for dynamic sizing, it is currently challenging to change the size of either component after installation. PV panels are commonly installed on rooftops, and often require a licensed technician to work with. In contrast, a battery is typically installed in a more accessible place, such as a garage, but is often enclosed in a protective casing. Changing the battery size requires re-imagining the battery from a monolithic structure to a set of smaller modules that may be added or removed safely and easily. The rest of our discussion in this section will focus on dynamic sizing of batteries.

Battery products such as Tesla's Powerwall [161] and Enphase's AC battery [40] already offer modular batteries that may be scaled up according to the needs of the system owner. However, a dynamic battery size implies the ability to both increase and decrease the size as necessary. Apart from the technical aspects regarding the modification of system hardware, there is also a need for a marketplace to support the buying and selling of new and used battery modules. To properly valuate used battery modules, there has to be a way to reliably determine the level of degradation and remaining battery service life, for which several methods have been developed [188, 170]. Support for a heterogeneous battery system, where combined battery modules may have different manufacturers, levels of degradation, and even cell chemistry, is a current area of study [15, 173, 118, 130] and would increase the scope of a used battery marketplace. Such a battery market, where there is no trusted third party, could be built as a trusted peer-to-peer economy based on

a blockchain with verifiable smart contracts executed on battery systems to estimate their remaining value. The participation of battery systems on the market may be completely autonomous, with each system periodically using collected data for computation of its own optimal size, followed by an arrangement of the necessary transactions on the marketplace, with the final decision made and realized by the system owner.

The idea of a truly dynamic battery size supported by a market is admittedly progressive. Nevertheless, as the cost of battery systems is expected to remain relatively high, and as it becomes ever more critical to replace conventional carbon-based power generation with cleaner alternatives, dynamic battery systems may add value, thus facilitating the widespread adoption of PV-storage systems.

**Chapter 4 Summary**
We study the problem of computing a robust PV-storage system sizing that meets the LOLP or EUE performance objectives while minimizing the system cost. We develop a novel framework for computing a robust system sizing based on data-driven system simulation and probability bounds. In contrast with other existing work, we relax the assumption that the available historical data which characterizes the system's operating environment is perfectly representative of the future and instead use a weaker assumption that the dynamics of both the past and future operating environments are characterized by the same probability distribution. Our approach is compared to a stochastic network calculus approach on a large dataset of household PV generation and load measurements.

# Chapter 5

# Optimization of System Operation

After carefully sizing a PV-storage system, it has to be operated effectively. To maximize system value, effective operating decisions should be made in a timely manner and minimize the cost of providing power to the user. This cost may be offset by using local PV generation, by selling PV generation to the grid, and by taking advantage of the grid pricing scheme.

The system's operating flexibility comes from the battery, and the main control variables are the rates of battery charging and discharging. The choice of these variables depends on the operating environment, including the grid price, locally generated power, and load, and possibly forecasts for these values. Different combinations of these factors may require different operating strategies.

In this chapter, we review existing approaches to system operation in Section 5.1, and provide a mathematical model of the system under consideration in Section 5.2. We then study three methods for system operation: rule-based algorithms in Section 5.3, model predictive control in Section 5.4, and neural networks in Section 5.5.

The contributions of this chapter are as follows. We develop three rule-based algorithms for the control of PV-storage systems, based on insights from optimal operation. We also propose and evaluate a mixed-timescale model predictive control approach which can achieve near-optimal performance even with a limited computational budget. Finally, we develop a self-learning controller based on neural networks that learns how to effectively operate a PV-storage system post-deployment, and effectively adapts its strategy to changes in the operating environment.

Materials in this chapter have been published in [79] and [77].

Table 5.1: Operation Notation

| *Name* | *Description* |
|---|---|
| $\beta$ | Exponent coefficient of exponential MPC timescale |
| $\delta$ | Number of low-resolution time slots in two-stage MPC timescale |
| $\Gamma$ | Peak-demand pricing scheme threshold |
| $P_g$ | Power purchased from the grid |
| $P_{over}$ | Amount of power exceeding peak-demand threshold |
| $P_{sell}$ | Power sold to the grid |
| $p_1$ | Price during ToU high-price interval (day) |
| $p_2$ | Price during ToU low-price interval (night) |
| $\pi_b$ | Base price for peak-demand pricing scheme |
| $\pi_d$ | Demand price for peak-demand pricing scheme |
| $\pi_g(t)$ | Price of buying power from the grid in time slot $t$ |
| $\pi_{sell}(t)$ | Price for selling power to the grid in time slot $t$ |
| $T_b$ | Time slot budget for MPC |
| $X_j$ | Energy content of $j$th high-price interval |
| $Y$ | Energy content reserved for offsetting peak consumption |
| $\zeta$ | Number of minute-long time slots in two-stage MPC timescale |

## 5.1 Review of System Operation Approaches

System operation approaches are evaluated on many characteristics. The overall performance of the approach, i.e., the degree to which the objectives of the operation are satisfied, is the most important characteristic. In the context of energy system control, we take a closer look at those techniques that exhibit good performance and have reasonable online complexity, as discussed below. Techniques such as model predictive control, rule-based algorithms, and neural networks have all exhibited good performance in the literature. These approaches differ in three main aspects:

1. **Online computational complexity:** the amount of computation required to compute a control action.

2. **Offline complexity:** the amount of human effort and computation required to configure the controller prior to its deployment in an online environment.

3. **Interpretability:** the degree to which a system owner is expected to correctly interpret the relationship between the input to the controller and the decisions made

by the controller. We use the relative terms "low", "medium", and "high" in our assessment of a technique's interpretability.

The ideal control approach would have low online computational complexity, low offline complexity, and high interpretability. The three aforementioned approaches to system operation offer different trade-offs in relation to these aspects; we discuss them below.

### 5.1.1 Model Predictive Control

Model predictive control (MPC) is a method for online system control that is based on a dynamic system model and a predictor of future input to the controller [138]. To compute an action at a given time, a control optimization problem is solved to determine the control actions that optimize the objective function over a finite time horizon. The first control action (corresponding to the current time) is implemented by the controller, and the process is repeated to compute the next control action with a shifted optimization horizon. MPC is commonly used for industrial process control [184], although there have been applications of MPC in a variety of other areas, including the optimization of energy efficiency [140] and HVAC systems [66] in buildings, as well as controlling battery and wind energy systems [178].

MPC approaches may be divided into linear and non-linear classes, which differ in the complexity of the system model. Linear MPC is the most prevalent [136] since, in many cases, MPC provides effective system control even when non-linear system behaviours are approximated by linear models. The difficulty with non-linear MPC lies in the computational burden of solving non-linear optimization problems, which is especially important in an online control setting [184].

MPC has been studied as a control method for energy systems with storage. For example, Teleke et al. [159] consider a system consisting of a wind farm and energy storage that is used to make wind power more dispatchable. One of the observations in this work is the system performance with respect to the length of the time horizon considered for optimization; horizon values ranging from 100 seconds to 30 minutes were tested, showing significant performance benefits from having longer – and hence, more computationally expensive – horizons.

Khalid et al. [84, 86] also consider a wind farm with energy storage but with different objectives. In [84], energy storage is controlled using an MPC approach to reduce the ramping of system power output, and in [86] the system is operated to maximize revenue from selling energy in an electricity market with dynamic pricing. In both cases, it is shown

that the MPC can effectively utilize the battery to meet the objectives of the applications. The optimization horizon in [86] is restricted to 15 minutes to limit the computational burden.

MPC has also been studied in the context of microgrid operation, which is influenced by many stochastic processes. For example, in [187], a Monte Carlo sampling technique is used to generate a set of forecasts for the operating environment, including load, electricity price, and PV and wind generation, which are used as input to a stochastic optimization problem. This led to improved performance over a standard MPC approach, at the cost of increasing the computational cost by a factor of 10. Another example [126] involves splitting the optimization step into two stages[1]: an initial stage involving a linearized optimization problem to determine an initial operating plan over a relatively long time horizon, and a secondary stage involving a non-linear model to calculate the power flows over a relatively short horizon to implement the initial plan.

MPC has a relatively high computational cost, since an optimization problem is solved every time a new action is computed. The dynamics of a PV-storage system suggest that a new action would be required at least once per minute, and the controller hardware typically has very limited computational capacity to minimize costs as well as overhead power consumption. There are many techniques for reducing the computational cost of MPC, such as hierarchical decomposition of the control space [142] and non-uniform hold constraints [106], but even these techniques may not sufficiently reduce the computation time if a non-convex system model is needed.

Compared to the other two techniques discussed in this section, MPC has a relatively low configuration cost. It requires a system model, prediction model for the operating environment, the objective of the system expressed as a mathematical function of the control actions, and a generic optimization problem solver such as CPLEX or Gurobi.

Finally, the actions computed by MPC are completely predetermined by the parameters and inputs of the model, and therefore interpretable. The degree of interpretability is a function of the system owner's familiarity with the system, system model, and prediction models. We assess the interpretability of MPC as medium to high.

### 5.1.2 Rule-based control

The rule-based control approach involves creating a set of rules which determine control actions depending on system state. Rule-based control typically requires very low online

[1]This two-stage approach is different from the two-stage time horizon discussed in Section 5.4

computational effort. However, coming up with rules that result in good system performance is not always a trivial task. Rules are specified in an explicit "white box" form such as an algorithm or knowledge-based system [155], making this approach highly interpretable. The creation of explicit rules is often done by domain experts.

It is common for rule-based strategies to be very simple, with a focus on meeting system constraints. In [169], a very simple algorithm is used to control an off-grid system composed of a wind turbine, PV panels, hydrogen fuel cell, and household loads. The algorithm chooses one out of two predetermined control actions depending on whether or not there is excess renewable power generation. A similar system was considered in [7], where the control of the system included a heuristic that kept the charge of the battery at high levels to reduce the problems in power balance caused by intermittent renewable sources.

A simple operating strategy that is commonly used in PV-storage systems is to charge the battery if there is an excess in PV energy ($P_S(t) > P_L(t)$), and discharge the battery if PV is unable to fully meet the load ($P_S(t) < P_L(t)$). The pseudocode for this strategy is given in Algorithm 1. This strategy maximizes the consumption of PV generation [26], and may be used as a simple benchmark. Note that this is a myopic strategy, i.e., it does not use predictions for PV generation and load, and it does not use any grid pricing scheme parameters.

**Algorithm 1** Pseudocode for simple control strategy.

**if** $P_S(t) > P_L(t)$ **then**
  Charge the battery as much as possible
  Sell the excess solar
**else**
  Discharge the battery to make up for the difference between $P_S(t)$ and $P_L(t)$
  Purchase power from the grid if the battery power was not enough
**end if**

In a general dynamic pricing scenario, it has been shown that an optimal operation policy can be based on energy thresholds [58]. More concretely, the charge/discharge policy may be described through a relationship between the energy content in the battery, some optimal energy thresholds, and the state of the system.

While rule-based approaches have shown to be effective, requiring a domain expert to tweak the battery control rules to adapt them to changes in the environment is not practical. Furthermore, there are a large number of grid pricing schemes; each utility company offers several options to its residential customers and an even larger variety of

options to commercial and industrial customers. Each pricing scheme may require its own set of tailored control rules. The offline complexity of this approach comes from the effort required for rule design.

### 5.1.3 Neural Networks

Artificial neural networks have been studied in the context of controlling energy systems such as hybrid-electric vehicles and wind farms with energy storage [108, 23, 117]. In [117], a neural network approach to controlling a supercapacitor in a hybrid-electric vehicle showed substantial improvement in energy efficiency compared to other approaches. In [23], a neural network is shown to decrease the cost of operation compared to simple rule-based control and fuzzy logic control. In both studies, the neural network was trained using datasets of "good" control actions which were obtained by using offline optimization methods.

Deep learning methods have shown promise in complex control problems. Deep neural networks have been trained to successfully perform physical tasks such as walking and car driving in a simulated environment [104], play video games [115, 116], and replace PID control for a DC motor [31]. These problems have high-dimensional features (input) that are mapped to a comparatively small number of control variables (output); energy storage control has the same challenge, where a large number of environment variables can be used to decide the battery's charging or discharging rate. Applying deep neural networks to energy systems has been studied in [45] and [110], and, recently, the concept of supporting and even replacing human effort via artificial neural networks in the management of power systems has been proposed [164].

The online computational cost of using a deep neural network is very low, making it more suitable for small-scale residential systems compared to MPC. In [186] and [62], the high computational cost of using MPC was handled by using neural networks to encode an online MPC policy to control drones and other autonomous robots. The shortcomings of this approach are the cost of training the network, and the need for large amounts of data to learn from.

Neural networks are regarded as a "black box" method. The interpretability of a deep neural network is notoriously low.

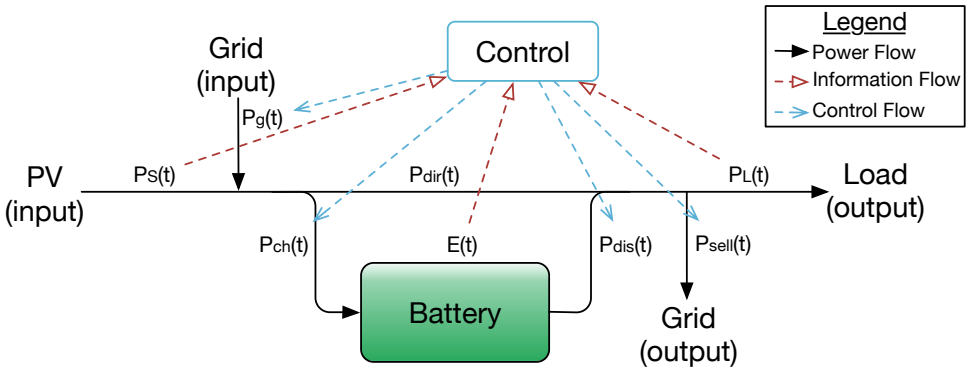


Figure 5.1: PV-storage system diagram for operation

## 5.2 Mathematical Model

We consider a PV-storage installation in a building consisting of:

- a set of PV panels along with associated inverters and power electronics; the PV output is assumed large enough to significantly reduce grid use;
- a Li-ion battery that can either store or discharge energy;
- a bi-directional grid connection that allows load to be met from the grid, the battery to be charged from the grid, and for power to be sold to the grid;
- a control component that operates the system in real time.

Figure 5.1 illustrates our system.

The PV-storage system is assumed to be owned by and operated in the interest of a home or small-business owner. The grid charges the owner for any electricity that is supplied, with the price depending on the pricing scheme being considered. The grid also purchases any electricity that the system chooses to sell.

In this section, we formulate the mathematical model of the system and implement the model as a mathematical program. The programs are used to compute the optimal operation of the system, which provides a benchmark for evaluation of the operating strategies developed in later sections.

### 5.2.1 Formulation of the model

We use a discrete-time model, with $T_u$ being the length of each time slot. For simplicity, we define $t$ to represent the time interval $[t \times T_u, (t+1) \times T_u)$, and use the phrase 'at time $t$' meaning 'during the time interval $[t \times T_u, (t+1) \times T_u)$'. All the model variables are assumed to be constant within a time slot, and the total number of time slots under consideration is denoted $T_h$. The available output power of the PV panels at time $t$ is equal to $P_S(t)$. The PV output power may be used to meet the load directly, to charge the battery provided that there is room, or to be sold to the grid. The grid may also be used to meet the load and/or charge the battery, with $P_g(t)$ being the power bought from the grid at time $t$. The load at time $t$ is equal to $P_L(t)$, and can be met using a combination of PV, grid, and battery output. The power sold by the system is denoted as $P_{sell}(t)$. $P_{dir}(t)$ is the power that flows directly from PV and grid sources to meet the load or be sold to the grid at time $t$. $P_c(t)$ is the power from the sources used to charge the battery, and $P_d(t)$ is the power from the battery used to serve the load at time $t$.

There are some natural constraints on the flows of power through the system. We require the load to be met at all times, and the energy to be sold either form the battery or from grid[2] and PV sources. Thus

$$P_{dir}(t) + P_d(t) = P_L(t) + P_{sell}(t), \forall t \in [1, T_h]. \tag{5.1}$$

In addition, the sum of $P_c(t)$ and $P_{dir}(t)$ must not exceed the input power of the system, therefore

$$0 \leq P_c(t) + P_{dir}(t) \leq P_S(t) + P_g(t), \forall t \in [1, T_h]. \tag{5.2}$$

We consider the battery to be a Li-Ion battery modelled using the C/L/C model described in Section 3.4.

We denote $\pi_{sell}(t)$ to be the cost per unit of energy sold to the grid at time $t$. We consider the electricity payment to be the sum of outgoing grid payments minus the sum of incoming payments that we receive from the grid for selling energy through either a feed-in tariff or value-of-solar tariff structure (as described in Section 2.1.4).

We consider two pricing schemes:
**Differential Pricing (also called real-time pricing):** Let $\pi_g(t)$ be the cost per unit of

[2]The system does not prohibit the immediate selling of power bought from the grid. In our work, we make the simplifying assumption that the pricing structure prevents this from being a profitable operation, i.e., the cost of buying from the grid is assumed to always be greater than the cost of selling to the grid.

grid energy at time $t$. The electricity payment over the time horizon is therefore

$$\sum_{t=1}^{T_h}(\pi_g(t)P_g(t) - \pi_{sell}(t)P_{sell}(t))T_u, \tag{5.3}$$

**Peak-demand Pricing:** The price of power from the grid increases if the system exceeds the grid purchasing power threshold $\Gamma$. The grid charges $\pi_b$ for every unit of energy purchased, with an additional $\pi_d$ demand price for the power demand that exceeds $\Gamma$. We denote $P_{over}(t)$ to be the purchased power that exceeds $\Gamma$ at time $t$:

$$P_{over}(t) = max(0, P_g(t) - \Gamma). \tag{5.4}$$

The electricity payment over the time horizon is therefore

$$\sum_{t=1}^{T_h}(\pi_b P_g(t) + \pi_d P_{over}(t) - \pi_{sell}(t)P_{sell}(t))T_u. \tag{5.5}$$

### 5.2.2 Offline Optimal Benchmark

The optimal offline operational strategy (i.e., the oracle) is based on the assumption that our system has complete knowledge of load and PV generation over the entire time horizon. The oracle is used to obtain a benchmark against which realistic (online) operating strategies, which do not have complete knowledge of the future, are evaluated.

Combining the system constraints and grid payment function, we formulate the problem of minimizing the electricity payment of the system owner for the two pricing schemes as linear programs (LP).

**Differential Pricing:** Given solar and load datasets $(P_S(t))$, $(P_L(t))$, the C/L/L battery model parameters $n$, $u_1$, $u_2$, $v_1$, $v_2$, $\alpha_c$, $\alpha_d$, $\eta_c$, $\eta_d$, $U$, the pricing data $(\pi_g(t))$, $(\pi_{sell}(t))$, and the time constants $T_u$, $T_h$:

$$\min_{\substack{P_{dir}(t),P_c(t),P_d(t),\\ P_g(t),P_{sell}(t)}} \sum_{t=1}^{T_h} (\pi_g(t)P_g(t) - \pi_{sell}(t)P_{sell}(t))T_u \tag{5.6}$$

subject to

$$0 \leq P_g(t), P_{dir}(t), P_{sell}(t) \quad \forall t \tag{5.7}$$

$$P_{dir}(t) + P_d(t) = P_L(t) + P_{sell}(t) \quad \forall t \tag{5.8}$$

$$E(0) = U \tag{5.9}$$

$$E(t) = E(t-1) + \eta_c P_c(t) T_u - \frac{P_d(t)}{\eta_d} T_u \quad \forall t \tag{5.10}$$

$$0 \leq P_c(t) + P_{dir}(t) \leq P_S(t) + P_g(t) \quad \forall t \tag{5.11}$$

$$0 \leq P_c(t) \leq \alpha_c n \quad \forall t \tag{5.12}$$

$$0 \leq P_d(t) \leq \alpha_d n \quad \forall t \tag{5.13}$$

$$P_c(t)P_d(t) = 0 \quad \forall t \tag{5.14}$$

$$u_1 P_d(t) + v_1 n \leq E(t) \leq u_2 P_c(t) + v_2 n \quad \forall t \tag{5.15}$$

**Peak-demand Pricing:** Given $(P_S(t))$, $(P_L(t))$, $B$, $u_1$, $u_2$, $v_1$, $v_2$, $\alpha_c$, $\alpha_d$, $\eta_c$, $\eta_d$, $U$, $\pi_b$, $\pi_d$, $\Gamma$, $(\pi_{sell}(t))$ $T_u$, $T_h$:

$$\min_{\substack{P_{dir}(t),P_c(t),\\ P_d(t),P_g(t),\\ P_{sell}(t),P_{over}(t),\\ I(t)}} \sum_{t=1}^{T_h} (\pi_b P_g(t) + \pi_d P_{over}(t) - \pi_{sell}(t)P_{sell}(t))T_u \tag{5.16}$$

subject to

Constraints (5.7-5.15)

$$0 \leq P_{over}(t) \quad \forall t \tag{5.17}$$

$$P_g(t) - \Gamma \leq P_{over}(t) \quad \forall t \tag{5.18}$$

Note that we do not define $P_{over}(t)$ as the equality in Eq. (5.4). Instead, the combination of two linear constraints (Constraints 5.17 and 5.18) is used with the same effect. Furthermore, the non-linear Constraint 5.14, which prevents simultaneous charging and discharging, can be relaxed; the optimal strategy will naturally avoid this action which causes energy losses due to battery inefficiencies, as proven in [48]. The resulting program is linear.

To better understand the effect of system operation on the system owner's payments to the grid, we solve the LPs using real solar and load data traces, and compare the minimized grid payments with what they would have been under the simple PV-storage operating strategy.

### Parameter values

We solve the LPs using the following values for the parameters. $P_S(t)$ and $P_L(t)$ are real traces from the Technology Center for Energy (TZE) in Ruhstorf, Germany for nine months, April through December in 2014, where measurements were recorded every five minutes ($T_u = 1/12$ hours). The maximum observed load is approximately 24 kW, while the maximum observed PV output is approximately 25 kW. The C/L/L battery model parameters reflect an NMC battery and are specified in Table 4.2. We assume that the battery has 50% charge at $t = 0$, i.e., $U = n \cdot v2 \cdot 0.5$. The capacity of the battery is set in the range from 10 to 100 kWh.

To imitate the fixed-rate contracts for buying solar power that are seen in practice, we set the selling price $\pi_{sell}(t)$ to be a constant value of \$0.12 per kWh, which is the approximate feed-in tariff in Germany as of August 2015.

**Differential Pricing:** The LP for differential pricing can compute the optimal strategy for any pricing scheme in which price differentials occur over time. For the numerical results, we consider a common type of differential pricing called time-of-use (ToU) pricing. We set a high price $p_1$ during the day [8am, 8pm) and low price $p_2$ during the night [8pm, 8am), resulting in a pricing that has the same pattern as the typical demand for user electricity consumption. We fix $p_2$ to be \$0.25 per kWh, which is reflective of typical spot-prices in Germany at night [109], and vary $p_1$ between \$0.375 and \$0.75.

**Peak-demand Pricing:** We set $\Gamma = 9.06$ kW, corresponding to the 80th percentile of load measurements over the nine month period. The base price $\pi_b$ is set to \$0.25, and the demand price $\pi_d$ is varied between \$1.50 and \$2.50.

Note that for both pricing schemes, the lowest price for purchasing power from the grid is higher than the payment obtained by selling power to the grid, effectively preventing grid arbitrage. Otherwise, the optimal strategy would be to buy as much power as possible from the grid and sell it for profit.

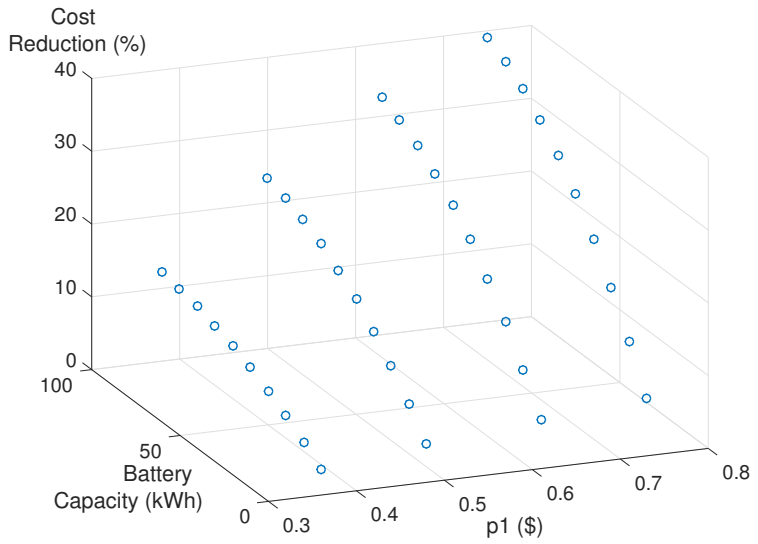


Figure 5.2: Cost reduction percentage from using the optimal strategy over the simple strategy under differential pricing

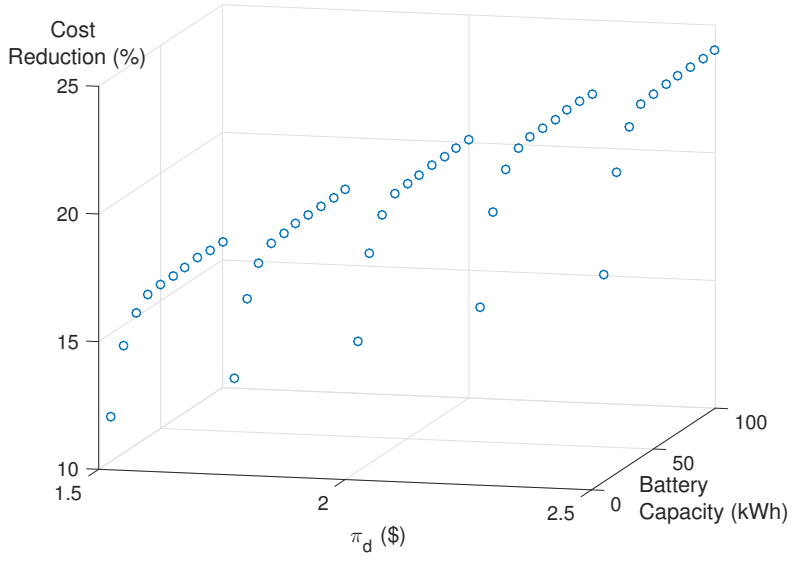


Figure 5.3: Cost reduction percentage from using the optimal strategy over the simple strategy under peak-demand pricing

#### Optimal Operation

Figures 5.2 and 5.3 show the cost savings of the optimal strategy in comparison with the simple myopic strategy (Algorithm 1) which is simulated on the system using the same inputs that were used to solve the LP. Significant cost savings are possible if the optimal strategy is used, which confirms the necessity of an operating strategy that is better than the simple myopic and price-agnostic strategy. These savings increase when the price ($p_2$ or $\pi_d$) increases, and when the size of the battery increases, although there is a critical battery size beyond which we see no increase in cost reduction for a finite time horizon.

To realize these gains, the oracle takes advantage of precise knowledge of the future, both in terms of load and PV generation. A deployment of an oracle-based approach on a real system would require accurate forecasts, and it is unclear how forecasting errors would affect the performance of the system. However, the optimal strategy provides valuable insights on which the design of a practical strategy may be based.

## 5.3 Rule-Based Algorithms

Our method is to study the oracle's operations, extract insights about how the oracle works, and then implement these insights in the form of an algorithm. In the development of algorithms for operating the PV-storage system, there are many similarities between differential and peak-demand pricing schemes. We avoid repetition of ideas by focusing on the peak-demand pricing scheme.

### 5.3.1 Insights on Optimal Operation

To understand the relationship between the operating environment and the operation of the oracle, we visualize the optimal charging and discharging patterns over time against the PV generation, load, battery state of charge (SoC), and grid price curves. Figure 5.4 is an example for three days with the peak-demand pricing scheme. The following observations highlight key points in the development of a control algorithm:

**Observation 1:** If there is an upcoming period where PV output is not sufficient for meeting the load, and the difference between them is greater than $\Gamma$, then the grid is used to cover only up to $\Gamma$ of the remaining load; the battery is discharged to cover the rest, if possible. If there is no such period in the near future, then the battery is discharged

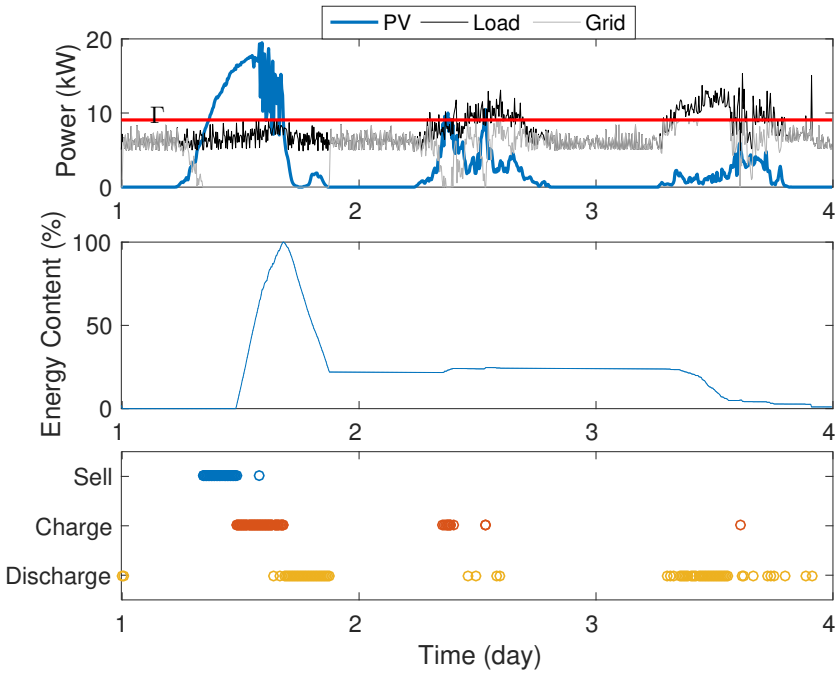


Figure 5.4: Visualization of the optimal strategy for peak-demand pricing over three days. PV, Load, and battery SoC are shown. Selling, charging, and discharging events (but not quantities) are also depicted.

whenever PV output is not sufficient for meeting the entire load, and the grid is used as a last resort.

**Observation 2:** The grid is used to charge the battery on some days if PV output in the near future is low and the battery does not have enough charge to prevent grid use from exceeding $\Gamma$ during that time. The charge level depends on the severity of the PV deficit; lower PV generation or greater load increases the amount of required charging. In the situation where the grid is used to charge the battery, the overall grid use does not exceed $\Gamma$.

**Observation 3:** The system sells power to the grid only when the battery is full and load has been met. Only PV power is sold, since it is not profitable to sell grid power because of the negative price difference. Discharging the battery to sell power is also sub-optimal due to efficiency losses incurred when the battery is charged again. This means that the selling price does not affect the optimal strategy, as long as it is lower than the lowest grid price.

In effect, the optimal strategy for peak-demand pricing balances two competing risks:

1. **The risk of grid use exceeding $\Gamma$.** We can avoid paying the high demand price by using the battery as a back-up to keep the grid usage below $\Gamma$. This risk is minimized if the battery is charged fully at all times except for when it is needed to prevent grid use from exceeding $\Gamma$.

2. **The risk of wasting PV output.** If PV generation exceeds the load and the battery is fully charged, then the excess generation is sold to the grid at a low price, as opposed to being stored and used to offset the purchase of grid power at a later time. This may be avoided by discharging the battery whenever PV generation is not sufficient for meeting the load, to make room for storing excess PV generation in the future. This risk is minimized by having the battery be empty as often as possible.

Similar visualizations of the optimal strategy may be used for the ToU pricing scheme. Having done so, we observe that the optimal strategy for ToU pricing balances two competing risks which are similar to those observed for peak-demand pricing. They are:

1. **The risk of using the grid during the day.** We can avoid paying higher prices during the day by charging the battery from the grid at night when prices are lower, and by using that energy during the day. This risk is minimized if the battery is fully charged by the beginning of the high-price period.

2. **The risk of wasting PV output during the day.** This can be avoided by discharging the battery at night (to meet the night load) to make room for PV generation which exceeds the load during the day. This risk is minimized by having the battery empty by the beginning of the day, i.e., by the beginning of the high-price period.

Using these insights, we develop practical control strategies which operate the system as closely as possible to what has been observed in the optimal operation.

### 5.3.2 Operating Strategies

We assume that, at time $t$, the system controller knows $P_S(t)$, $P_L(t)$, $E(t)$, and possibly some aggregate predictions for total PV generation and load in the near future, and is responsible for computing $P_g(t)$, $P_e$, $P_c(t)$, and $P_d(t)$. We further assume that the system controller is aware of the pricing scheme. Our goal is to develop an algorithm, i.e., a set of rules for the controller, that uses the limited available information to reduce the grid payments by the system owner. A practical strategy controls the operation of the system at any time while relying only on easily-obtained information.

The algorithms that we describe in this section rely primarily on myopic measurements although some of them also rely on a limited amount of predictions. They make feasible control decisions and have very low time complexity. Therefore, they can be easily implemented on a real system, though they do require expert input to design the algorithm. In the following, we describe an existing strategy [189] for two-period ToU pricing that makes decisions based on a parameter $X_j$, that refers to the amount of energy in the battery at the beginning of the $j$th high-price period. We also present a new strategy for peak-demand pricing which uses a single parameter $Y$ referring to the amount of energy in the battery maintained as a back-up, meant to avoid paying the demand price.

#### Strategy for ToU Pricing

To a certain degree, the algorithm given by Zhu et al. [189] takes into account the insights that we gained from studying the optimal strategy for ToU pricing. Their algorithm has two main cases, depending on the price period. During the high price period, the algorithm behaves similar to the simple strategy, charging only with excess solar and discharging whenever needed to avoid using the grid. During the low price period, the algorithm is effectively preparing to avoid using the grid in the next high price period. The battery is charged if its energy content is below $X_j$, and allowed to be discharged if it is above $X_j$.

We make the following modifications to the pseudo-code in [189] as Algorithm 2. We use the notation $[z]_+$ to represent $max(0, z)$. The modifications are:

**Modification 1:** The algorithm in [189] was not designed to consider the opportunity to sell power to the grid, and instead curtails the PV power if it cannot be used or stored. In the optimal strategy, only the excess solar power is sold, so we modify the algorithm to do this, i.e., $P_{sell}(t) = [P_S(t) - P_L(t) - P_c(t)]_+$.

**Modification 2:** The algorithm in [189] charges the battery as soon as possible using PV and grid sources once a decision is made for it to be charged, in order to have enough energy to meet the demands of the upcoming high-price period. We modify the algorithm to charge the battery at the latest possible time, i.e., right before the beginning of the high-price period. This is done to minimize grid use by giving the battery a chance to be charged with excess PV power before the high-price period begins.

**Algorithm 2** Storage operation at any time $t$, given $X_j$

**if** $t$ is in a low-price period **then**
  **if** $E(t) < X_j$ and $t_{jstart} - t \leq \frac{X_j - E(t)}{B\alpha_c\eta_c}$ **then**
    $P_c(t) = min\left(\frac{X_j - E(t)}{\eta_c T_u}, B\alpha_c\right)$,
    $P_d(t) = 0$,
    $P_g(t) = [P_L(t) - P_S(t) + P_c(t)]_+$
  **else**
    $P_c(t) = min\left([P_S(t) - P_L(t)]_+, B\alpha_c, \frac{a_2(P_c(t))^* - E(t)}{\eta_c T_u}\right)$,
    $P_d(t) = min\left(\frac{(E(t) - a_1(P_d(t)))\eta_d}{T_u}, \frac{(E(t) - X_j)\eta_d}{T_u}, B\alpha_d, [P_L(t) - P_S(t)]_+\right)$,
    $P_g(t) = [P_L(t) - P_S(t) - P_d(t)]_+$
  **end if**
**else if** $t$ is in a high-price period **then**
  $P_c(t) = min\left([P_S(t) - P_L(t)]_+, B\alpha_c, \frac{a_2(P_c(t))^* - E(t)}{\eta_c T_u}\right)$,
  $P_d(t) = min\left(\frac{(E(t) - a_1(P_d(t)))\eta_d}{T_u}, B\alpha_d, [P_L(t) - P_S(t)]_+\right)$,
  $P_g(t) = [P_L(t) - P_S(t) - P_d(t)]_+$,
**end if**
$P_{sell}(t) = [P_S(t) - P_L(t) - P_c(t)]_+$

* These equations are expressed compactly using a recursive definition. In practice, we solve the equations by searching for maximum values of $P_c(t)$ and $P_d(t)$ that satisfy the energy limit constraints.

The algorithm is critically dependent on the parameters $X_j$. A well-chosen $X_j$ value balances the risks described in Section 5.3.1. Zhu et al. propose that $X_j$ should be dynamic,

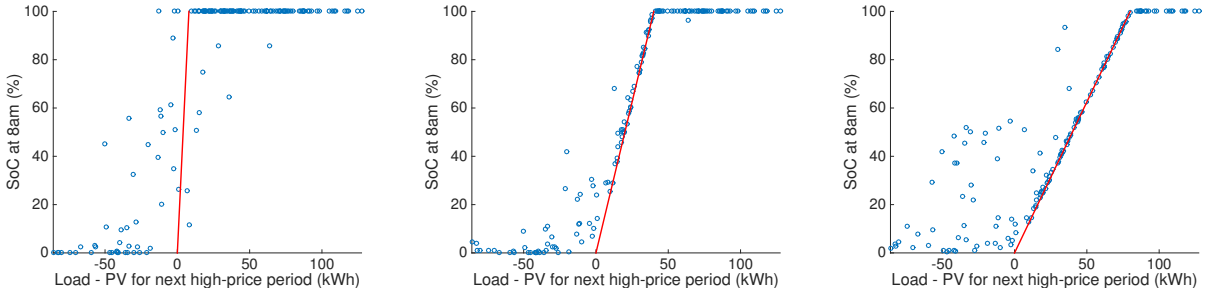


(a) $B = 10$ kWh (b) $B = 50$ kWh (c) $B = 100$ kWh

Figure 5.5: Optimal $X_j$ plotted against the day-ahead difference in load and PV generation. Each point represents one day's optimal $X_j$ value. The red line represents the choice of day-ahead $X_j$ as proposed by Zhu et al. [189].

with the controller predicting the PV generation and load for the high-price period of the upcoming day and setting $X_j$ to be the difference, i.e.:

$$X_j = \left[ \sum_{t=t_{jstart}}^{t_{jend}} ((P_L(t) - P_S(t))T_u \right]_{B \cdot v_1}^{B \cdot v_2} \tag{5.19}$$

where $[z]_y^x = max(x, min(y, z))$, and $t_{jstart}$ and $t_{jend}$ are the start and end times of the high-price period of day $j$. We will refer to this as the 'day-ahead $X$' strategy, and consider this to be a practical approach because the required predictions are made only for aggregate values, rather than for fine-grained load and PV generation curves. As it turns out, and as shown in Figure 5.5, the values of $X_j$ chosen using this method are very close to the energy content at the beginning of the high-price periods that we observe in the optimal strategy.

In the case where predictions for day-ahead PV and load are unavailable or unreliable, or if we wish to avoid predictions altogether, we simplify the problem of selecting the $X_j$'s by forcing $X_j = X$ for all $j$'s. We will refer to this simplified version of the algorithm as the 'static $X$' strategy. Note that, in practice, the controller would choose the value of $X$ based on data sets measured in the past at the premises and use the best value of $X$ for the given grid prices and the given battery size. The best $X$ for a given location, grid prices, and battery size may be seasonally dependent; we defer this analysis to future work. We evaluate the performance of the static and day-ahead $X$ strategies separately.

### Strategy for Peak-Demand Pricing

Our operating strategy for peak-demand pricing is developed on the base of our observations of the optimal strategy. Our strategy uses a static parameter $Y$, which is the amount of energy we attempt to have available as a backup when the load exceeds $\Gamma$. We will refer to this strategy as 'static $Y$'. The strategy is split into two operating modes, depending on whether the energy content of the battery is above or below $Y$ kWh. The mathematical expressions for defining each mode are more complex than those in Algorithm 2. For readability, we describe the static $Y$ strategy as a sequence of priorities for the flow of power through the system:

**Mode 1: if** $E(t) \leq Y$

1. PV power is used first to meet the load that exceeds $\Gamma$, then to charge the battery to $Y$. PV is then used to meet the rest of the load, and then to charge the battery if all the load has been met; remaining power is sold.
2. The battery is then discharged only if needed to meet the load that exceeds $\Gamma$.
3. The grid is used to meet the remaining load, and then to charge the battery to $Y$ as long as grid use does not exceed $\Gamma$.

**Mode 2: if** $E(t) > Y$

1. PV power is used to meet the load, then to charge the battery. The remaining power is sold.
2. The battery is discharged to meet any load that exceeds $\Gamma$. The battery is then further discharged, provided that the battery energy content does not fall below $Y$.
3. The grid is used to meet any remaining load.

Just like in the static $X$ strategy, in the static $Y$ strategy the $Y$ value does not require forecasting and may be chosen by training the algorithm on historical data. These static strategies are myopic, in the sense that they do not use predictions for PV generation and load. However, they do use the fact that past trends are useful for understanding what will work in the future, and this knowledge is incorporated with the $X$ and $Y$ parameters. In our numerical examples, we compare these strategies and observe their performance through a system simulation campaign.

### 5.3.3 Evaluation and Discussion

We simulate the PV-storage system using the proposed operating strategies, as well as the simple myopic strategy (Algorithm 1) for comparison. Our performance metric is the cost increase relative to the cost achieved by the optimal strategy. Given a strategy $S$, the relative cost increase is

$$\frac{Cost(S) - Cost^*}{Cost^*} \tag{5.20}$$

where $Cost^*$ is the cost given by the oracle. The battery size varies between 10 and 100kWh, which is a realistic range for the PV and load traces used in the simulation. We look at differential (ToU) and peak-demand pricing with the same prices as in Section 5.2.2. For determining the $X$ and $Y$ parameters, we incorporate a training and test set. We use data from April, June, August, October, and December to learn the best static $X$ and $Y$, and use them to test on May, July, September, and November data.

#### Differential Pricing Results

Figure 5.6 shows the cost reduction potential for each strategy. The day-ahead $X$ strategy performs near-optimally (within 4%) for all the tested battery sizes and grid price values as long as we are able to perfectly predict the aggregate PV and load for the next high-price period. The static $X$ strategy shows little cost increase (within 5%), and works best when the relative difference in day/night grid prices is large and/or battery size is small. A smaller battery benefits little from prediction, partly because there is less flexibility to decrease the electricity payments. Altogether, from these results, one might conclude that there is little need for forecasting.

The effects of varying the selling price were tested and were found to be minimal. This may be explained by noting that, in the tested scenarios, very little energy is sold by all the considered strategies. Indeed, our strategies are designed to use or store as much PV power as possible, with selling used as a last resort.

To test the robustness of the day-ahead strategy to prediction errors, we introduce errors into our simulation. Prediction errors, as a percentage of the actual value, are introduced by sampling a normal distribution with a mean of zero and applying it to the actual aggregate PV generation calculated for each day's high price period. According to the results described in [160], this is a reasonable distribution of prediction error. Similarly, we apply errors to the aggregate load predictions. We set the standard deviation $\sigma$ of the error distribution to be equal for both load and PV generation, in the range between 30% and

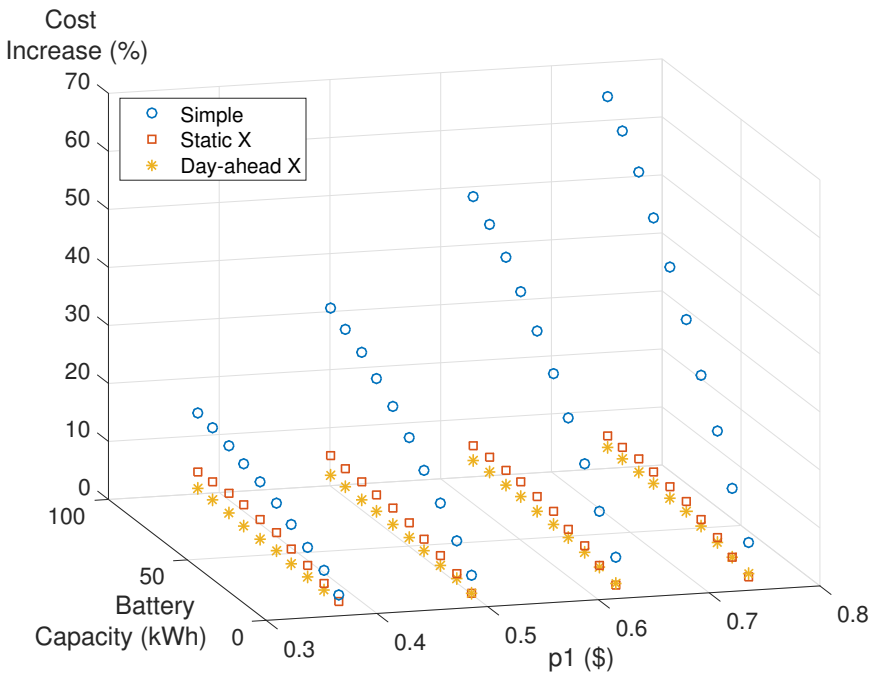


Figure 5.6: Cost increase across different battery sizes and $p_1$ values for the three strategies for ToU pricing.

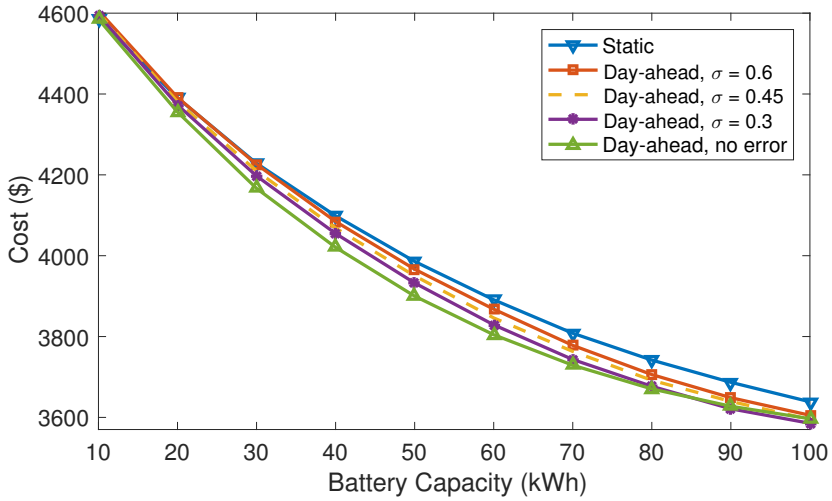


Figure 5.7: Effect of prediction error on the day-ahead strategy performance. $p_1 = 0.50, p_2 = 0.25$.

60%. We simulate the system with these prediction error distributions 20 times to obtain an average grid cost of the system over a five-month period for each $\sigma$. Figure 5.7 shows the effect of the error on the day-ahead strategy. These results support the conclusion that the day-ahead strategy is only marginally better than our static strategy. With substantial prediction errors ($\sigma = 60\%$) for both load and PV, the average performance is similar to the static strategy across most of the battery sizes and price ratios tested.

### Peak-demand Pricing

For peak-demand pricing, a second existing operational strategy, in addition to what we refer to as the simple strategy, is currently in commercial use, for example by Stem [154], a company that provides energy storage and management services. This strategy is designed for peak-demand pricing but does not consider a PV power source. It aims to use storage solely to prevent grid use from exceeding the demand threshold. We refer to this strategy as the 'back-up' strategy (because the battery is used entirely as a back-up power source), and model it as always being in Mode 1 of our static $Y$ strategy, i.e., set $Y = n \cdot v_2$, which is the upper limit on the battery energy content.

In Figure 5.8, we compare the relative cost increase of the simple, static $Y$, and back-up strategies. Our static $Y$ strategy exhibits near-optimal performance. The back-up strategy

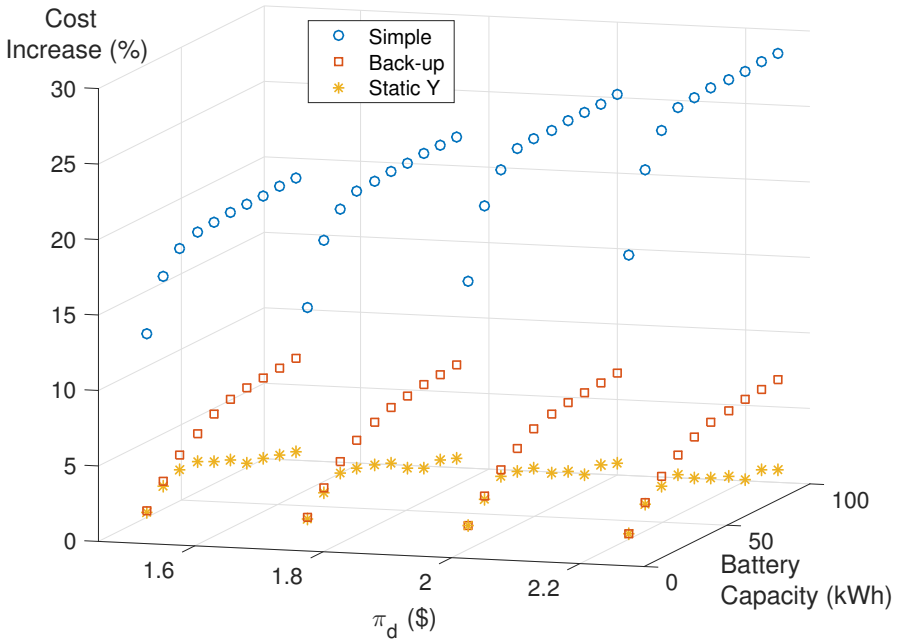


Figure 5.8: Cost increase across different battery sizes and $\pi_d$ values with peak-demand pricing of simple, static $Y$, and back-up strategies. $\pi_b = \$0.25$

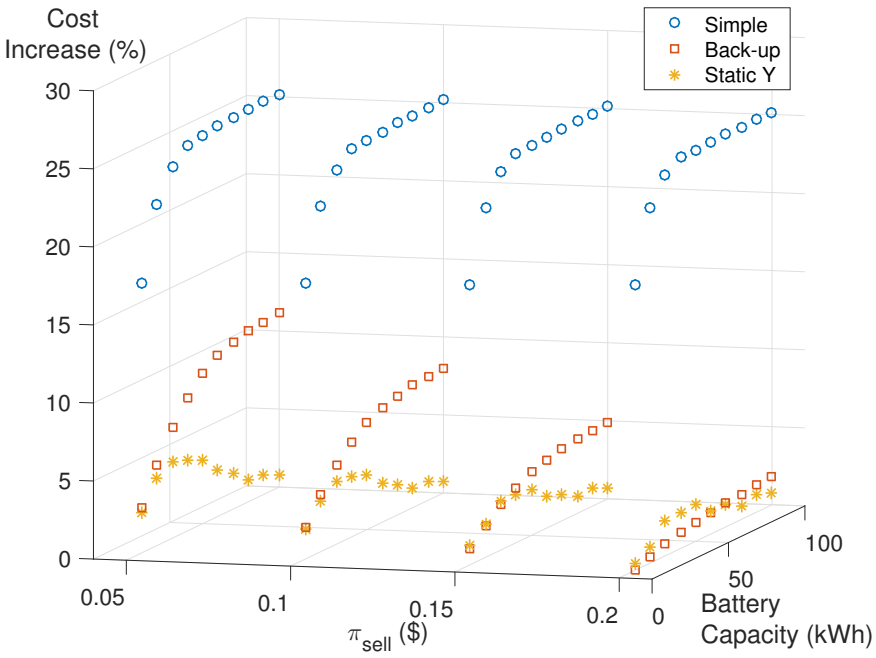


Figure 5.9: Cost increase with varying selling price and battery size. $\pi_b = \$0.25$ and $\pi_d = \$2.00$.

also exhibits a noticeable improvement over the simple strategy, though not as good as the static $Y$ strategy. This is not surprising, since the back-up strategy is a special case of the static $Y$ strategy. The back-up strategy performance degrades when the selling price is lower, as shown in Figure 5.9. This is because the battery is almost always kept full, and is unable to store excess PV generation which is consequently sold to the grid. These results suggest once again that forecasting might not be necessary, as the potential decrease in grid payments is less than 5%.

### Summary

Our work provides two key insights.

Firstly, the role of storage in this system is two-fold. On the one hand, it is used to store energy during low-price periods to reduce grid use during high-price periods or periods of high demand. On the other hand, it is used to store excess solar production (over demand) to reduce grid use during times of low (or zero) solar production. These two roles are sometimes in opposition: under some circumstances, it is better to have the battery fully charged, whereas, under other circumstances, it is best to have it fully discharged. This is why our heuristic strategies focus on $X_j$ and $Y$ as the critical tuning parameters, in an effort to find the best balance between the two uses.

Secondly, we propose rule-based strategies that do not rely on forecasting. In the ToU pricing case, we found that adding some aggregate forecasting may result in slightly better performance, and it is questionable if this benefit is worth the cost of the additional complexity, given that our static strategies already exhibit near-optimal performance when trained on representative data. Our approach to designing operating strategies, with a focus on avoiding complexity (i.e., forecasting, online optimization), shows promising results.

While our results with rule-based algorithms look very good, they depend on the assumption that past trends represent the future. Adapting our prediction-less algorithms to different pricing values may be achieved by recomputing static $X$ and $Y$ parameters, but a change to the pricing scheme structure would require a new set of rules designed by an expert. In general, rule-based algorithms are tailored to a specific deployment scenario and need to be re-designed if the scenario changes, for example to changes in the grid pricing scheme, load, and other events that significantly alter the operating environment. The rest of this chapter focuses on operation approaches that are better suited to adapting to a dynamic operating environment.

## 5.4 Mixed-Timescale Model Predictive Control

In the context of small-scale system control, such as a residential PV-storage system, the primary limiting factor of deploying MPC is the online computational cost of solving the underlying optimization problem to decide an action. The computation cost is dependent primarily on the number of variables in the optimization problem. In turn, the number of variables depends on the length of the optimization horizon. In Section 4.2.3 we demonstrated that the best-known polynomial time algorithm for solving linear programs has a complexity whose dominant term is approximately $T^{5.5}$, where $T$ is the number of time slots in the optimization horizon; this shows that every additional time slot in the optimization horizon comes at a high computational cost.

Given the computational resources available to the system controller, it is possible to determine a time slot *budget*: the maximum number of time slots, or equivalently, the optimization horizon with which the controller could run MPC quickly enough to compute timely control actions.

To determine the required frequency with which actions need to be computed, one could look at the timescales on which the relevant system parameters are changing. For a PV-storage system, the relative parameters that need to be considered by the controller are the incoming PV generation, load, and pricing scheme; of these, the PV generation and load typically vary more frequently, and their dynamics can be reasonably captured with a minutely timescale. Based on this observation, we may assume that actions should also be computed on a minutely timescale.

In experiments on hardware that are similar to what could be reasonably expected to be available to the system controller (Raspberry Pi), it has been measured that solving an MPC whose underlying optimization problem is an integer linear program (ILP) takes approximately 54 seconds when the optimization horizon has 24 time slots [135]. This experiment can be used to assess the time slot budget's order of magnitude. The budget would be larger if the underlying optimization problem is linear, possibly on the order of hundreds of time slots. The notion of a budget persists regardless of the complexity of the optimization problem.

Let us consider an example where the budget is 30 time slots. Given that the dynamics of the system are captured on a minutely timescale, one could set the optimization horizon to 30 minutes, with one minute per time slot. This approach may result in poor performance of the controller, since the MPC would not take into account any events that may be predicted to occur more than 30 minutes later. Alternatively, one could stretch some time slots beyond a minutely timescale, for example to five minutes per time slot.

The resulting optimization horizon would be longer and cover the future events that the controller would benefit from considering, at the cost of lower resolution of the predictions over the horizon. This consideration leads to the notion of *non-uniform hold constraints* [106, 163], in which the MPC's optimization horizon has a non-uniform timescale. This non-uniform timescale is typically set up with high-resolution time slots near the beginning of the horizon, and lower-resolution time slots toward the end of the horizon.

In this work, we explore two different ways to vary the timescale to improve the performance of an MPC while satisfying a given time slot budget. In the first approach, the timescale is split into a mix of predetermined fixed-length high-resolution time slots at the beginning of the horizon and low-resolution time slots toward the end of the horizon, which requires the configuration of two parameters: the number of high-resolution time slots, and the length of the remaining time slots. In the second approach, the length of each timeslot is determined using an exponential function and requires the configuration of a single parameter: the coefficient of the exponent. Using real data, we find that these parameters need only be configured once on the base of up to a single month of test data. The comparison of these two approaches shows that the exponential approach is more efficient: provides better results and has a simpler configuration.

### 5.4.1 Time Horizon Considerations

The simplest way to set up an MPC's optimization horizon is to set its length to be equal to the time slot budget; for a process with dynamics on a minutely time scale, this means a 30 time slot budget corresponds to an optimization horizon with a length of 30 minutes. We refer to this method as 'simple MPC'. When the budget is large, this approach works well. In the more common case where the budget is relatively small, this approach may result in poor performance, as discussed next.

In the context of residential PV-storage systems, we suppose that a 24 hour MPC horizon is sufficient for the achievement of quasi-optimal performance. We base this supposition on the following considerations. The time horizon must be at least as long as the typical cycle in the periodicity of the processes that have a significant effect on the operation of the system. We have three such processes:

- PV generation, which has a strong correlation with the time of the day and has a cycle of 24 hours.
- Electrical load, which in a residential home is dependent on the activities of the inhabitants, also exhibits cyclic behaviour repeating more or less every 24 hours.

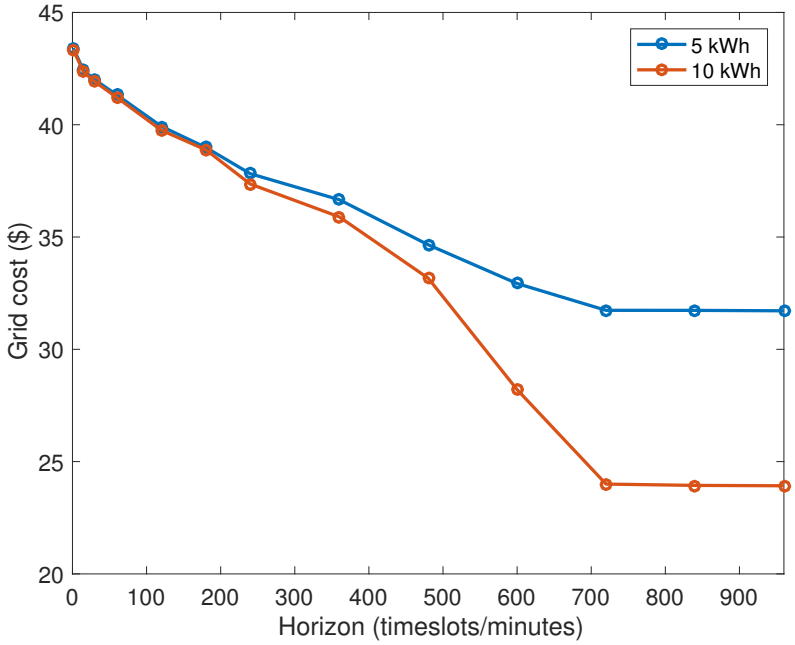


Figure 5.10: Grid payments with MPC operation over one month as a function of the optimization horizon, for ToU pricing with two battery sizes.

- Price of electricity, which in many cases is based on electricity demand, and thus conforms to a 24-hour cycle. Time-of-use pricing schemes are a typical example.

Even though some factors affecting these processes may have longer cycles, such as the season, we assume that the system may operate effectively without directly considering such long time scales[3] in the optimization horizon.

Experimentally, we have determined that time horizons shorter than 24 hours are also adequate for effective MPC performance. The simple MPC used in the experiments is formulated in Section 5.4.2, and used to control a PV-storage system in a simulation where the operating environment is represented by data traces measured from homes in Austin, Texas [1]. Figure 5.10 shows one month of grid payments under ToU pricing as a function of the length of the MPC's optimization horizon. We see that there are no further decreases in grid payments past a 720 minute (12 hour) horizon. This is consistent with

[3]Seasonality could and must be taken into account by the predictive models that are used by the MPC.

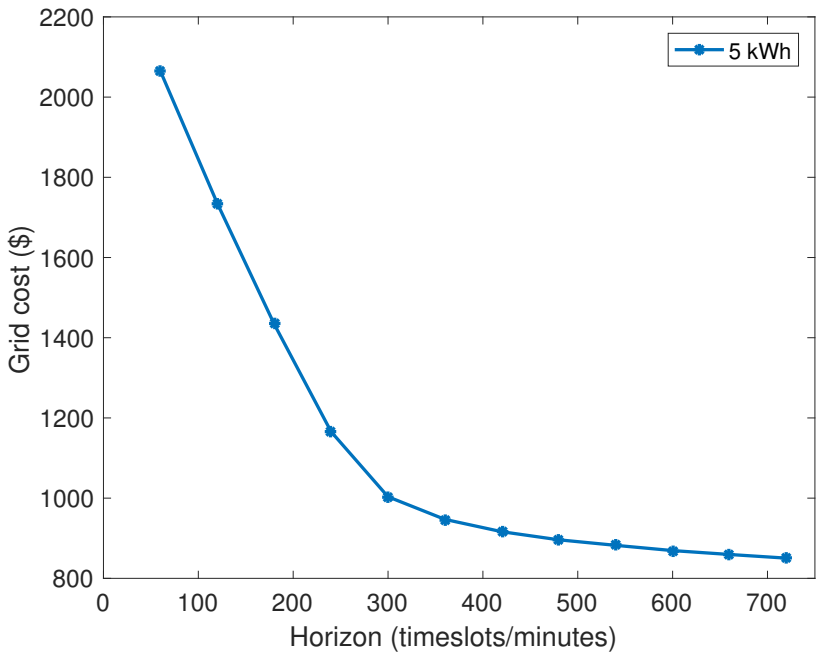


Figure 5.11: Grid payments with MPC operation over one year as a function of the optimization horizon, for peak-demand pricing.

the ToU pricing scheme used in this example, which has alternating 12 hour periods of low and high energy prices.

As a secondary observation, we note that the grid cost is almost identical for both 5 kWh and 10 kWh battery systems with a horizon of up to 200 minutes. This suggests that short horizons may eliminate the benefits of having a larger battery.

Figure 5.11 shows one year of grid payments under peak-demand pricing for a different home, as a function of the MPC horizon. We see that there is a knee-point in the curve at roughly 360 minutes (6 hours), after which the rate of improvement drops significantly.

The number of time slots in the budget may be lower than the number of time slots that we would prefer to fit into the desired optimization horizon. For example, Figure 5.11 shows that the optimization horizon should cover at least 360 minutes, whereas, according to the results published in [135], the computational budget is likely to allow only 20-30 time slots.

Due to the nature of the computations used by mathematical program solvers, the length of a time slot makes little or no difference to the computation time per slot. In other words, a long time slot contributes as much to the total computation time as a short time slot. Therefore, the budget depends on the number of slots and not on the resolution of each slot. Hence, a sufficiently long optimization horizon may be achieved by varying the length of time slots while staying within the computational budget of the controller.

One naive approach to fit a sufficiently long optimization horizon into a small time slot budget is to divide the horizon equally among the time slots. For example, consider a budget of 36 time slots; since 6 hours is a sufficiently long horizon for system operation via MPC for a peak-demand pricing scheme, we could split 6 hours into 36 time slots, each with a length of 10 minutes. The value of PV generation and load in each time slot would be the average of what is expected to occur in those 10 minutes. In doing so, we create a problem: the first time slot in the optimal solution, whose action is implemented by the real system, represents an aggregate of the load and PV generation we expect over the next 10 minutes, while the actual load and PV generation exhibit variation within those 10 minutes. As such, the action suggested by the MPC assumes a particular state of the system which may not be the actual state at that point in time. In this case, implementing the action requires some additional work to ensure that the power flow constraints of the system are not violated as a result of this discrepancy. Specifically, any deficit in the power flowing through the system, which may occur when the computed action is to charge the battery with PV generation that is not currently available, must be covered by purchasing power from the grid, and likewise, any surplus must be sold to the grid. The optimizer doesn't take these adjustments into account, and hence the resulting cost of operating the

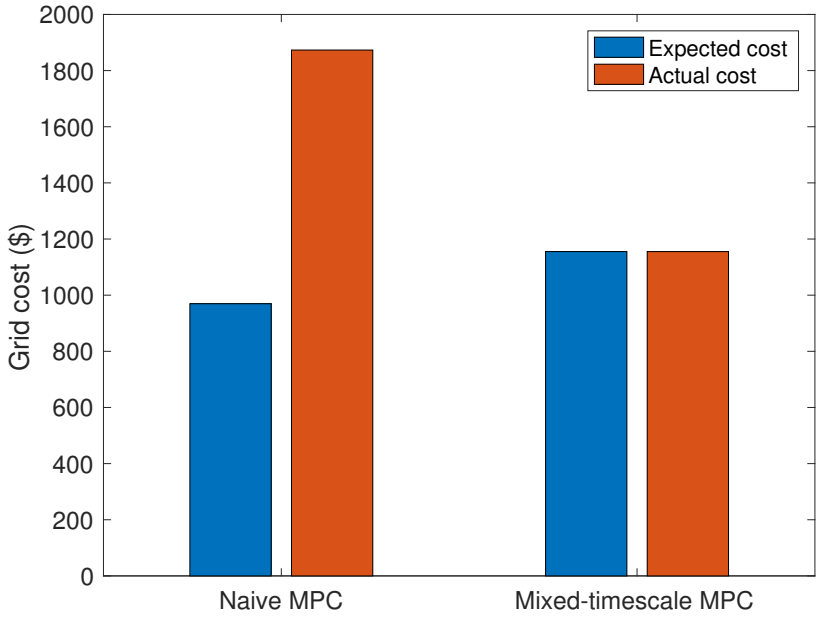


Figure 5.12: Cost comparison of what the MPC optimizer expects the grid cost to be, and the actual grid cost after balancing the power flowing through the system.

system may be very different from what is expected.

To get a sense of the difference in cost, we implement this naive MPC set-up in a simulation of a PV-storage system in a residential home over a period of one year, with the cost determined using the peak-demand grid pricing scheme. The resulting difference between the cost that the MPC thinks is being paid, and the actual cost after power balance is restored, is shown in Figure 5.12. In this example, the actual cost turns out to be approximately 1.9x larger than the expected cost.

The issue with the naive MPC is that it assumes a particular set of power flows representing the state of the system in the first time slot are feasible, when in reality this state may not be feasible and must be augmented by either buying or selling power. To rectify this problem, the first time slot in the MPC must represent the actual state of the system, i.e., it needs to be shorter. Returning to our example, by setting the very first time slot to be one minute long the remaining 35 time slots to be 10 minutes long, the resulting optimization horizon is still approximately 6 hours long. Implementing this in a simulation, we note that the actual grid cost is approximately 60% of the cost computed

with the naive MPC.

These experiments highlight two powerful ideas. First, the length of each time slot in the optimization horizon does not have to be equal. Second, the performance of the MPC depends significantly on the way the optimization horizon is split. In the rest of this section, we discuss ways in which the horizon may be effectively split into a mix of time slot lengths, resulting in what we refer to as "mixed-timescale MPC".

The results in Figures 5.10 and 5.11 are computed using time slots with a length of one minute. From the point of view of MPC's ability to satisfy the optimization objective, short time slots are preferable because they provide a lower degree of approximation and aggregation of the data reflecting the real processes. On the other hand, shorter time slots unavoidably lead to an increase in the computational cost.

By analogy with digital images, where smaller pixels mean higher resolution, in this section, short time slots are referred to as a *high-resolution* time slots. The shortest time slots that we consider are one minute long. This seems to be appropriate because the dynamic behaviour of the processes affecting system operation may be reasonably captured in minutely data traces. Accordingly, time slots with a duration greater than one minute are referred to as *low resolution* time slots.

The optimization horizon covers the current and expected future operating environment. In the context of prediction, at each point in time, the near future appears to be more definite than the distant future. Consequently, our ability to accurately and precisely predict near-future events is typically higher compared to distant events. For this reason, it makes sense to express the near future with high-resolution time slots, and distant future with low-resolution time slots.

Besides, the events in the near future have a relatively more pronounced effect on the present control action. For this reason as well, putting greater emphasis on the near future by expressing it with higher resolution is justified.

### 5.4.2 MPC Formulation

In this section, we formulate the MPC-based controllers for PV-storage system control for time-of-use and peak-demand grid pricing schemes. These formulations are similar to those presented in Section 5.2.1, and form the basis of an MPC for each pricing scheme. Forecasting models for PV generation and load are also an important part of the MPC, however, in this work, our study focuses on the optimizer.

Figure 5.1 serves as a guide for some of the notation and shows the label on each power flow. The power flowing through the system in each time slot $k$ is constant. Due to the conservation of power, the amount of PV generation and purchased grid power cannot exceed the sum of power used to charge the battery and the power that flows directly to meet the load or be sold ($P_{dir}$). Likewise, the amount of power being sold and demanded by the load cannot exceed the sum of $P_{dir}$ and the amount of power being discharged.

The battery has physical constraints that are described in our model. We use the C/L/C Li-ion battery model presented in Section 3.4 and briefly summarized here. The battery energy content at the end of time slot $k$ is denoted $E(k)$. The limits on battery energy content have a linear relationship with the charging/discharging power being applied, as described by the $u_1$, $u_2$, $v_1$, and $v_2$ parameters. The battery also has limits on the charging and discharging power, denoted $\alpha_c$ and $\alpha_d$ respectively. A fraction of power is lost when the battery is charged and discharged due to imperfect efficiency of power conversion, and the remaining fraction is denoted $\eta_c$ for charging and $\eta_d$ for discharging. The battery cannot be charged and discharged simultaneously, which is also enforced in our model.

The number of time slots in the budget is $T_b$, and the $k$th time slot is of length $T_u(k)$ hours, where $k$ is in the range $[1 \dots T_b]$. The values of $T_u(k)$ are among the inputs to the optimization problem. The optimization horizon is equal to the sum of all $T_u(k)$. We consider two timescale options for specifying $T_u(k)$: two-stage and exponential. The two-stage timescale specifies $T_u(k)$ as the following:

$$T_u(k) = \frac{1}{60} \qquad \forall k \in [1, \zeta] \tag{5.21}$$

$$T_u(k) = \delta \qquad \forall k \in [\zeta + 1, T_b], \tag{5.22}$$

where $\zeta$ represents the number of minute-length time slots, and $\delta$ represents the length of the rest of the time slots. In the special case where $\zeta = T_b$, this becomes a one-stage (simple) MPC.

The exponential timescale specifies $T_u(k)$ as the following:

$$T_u(k) \approx \frac{1}{60} e^{\beta(k-1)} \qquad \forall k \in [1, T_b]. \tag{5.23}$$

where $\approx$ indicates rounding to the nearest minute. As shown in Eqs. (5.21)–(5.23), the two-stage timescale has two parameters, $\zeta$ and $\delta$, while the exponential timescale has only one parameter $\beta$.

The formulations of the mixed-timescale MPC optimization problems for time-of-use and peak-demand pricing schemes are given in the following. Note that these specific

problems are linear, and may be implemented on inexpensive hardware such as a Raspberry Pi with a large time slot budget. Studying linear problems lets us simulate mixed-timescale MPC performance over long periods within a reasonable time frame.

### Time-of-use pricing

We denote $\pi_g(k)$ and $\pi_{sell}(k)$ to be the price for buying and selling energy in time slot $k$, respectively.

Combining our system constraints, we formulate the problem of minimizing the electricity payment by the system owner under time-of-use pricing as an integer linear program at time $t$:

Given the time slot budget $T_b$, the solar and load predictions: $[P_S(1), \ldots, P_S(T_b)]$, $[P_L(1), \ldots, P_L(T_b)]$; the battery parameters: $\alpha_c$, $\alpha_d$, $\eta_c$, $\eta_d$, $u_1$, $u_2$, $v_1$, $v_2$; the current energy content $U$; the prices: $[\pi_g(1), \ldots, \pi_g(T_b)]$, $[\pi_{sell}(1), \ldots \pi_{sell}(T_b))$, as well as the time slot lengths $[T_u(1), \ldots, T_u(T_b)]$:

$$\min_{\substack{P_{dir}(k),P_c(k),P_d(k),\\ P_g(k),P_{sell}(k),I(k)}} \sum_{k=1}^{T_b} (\pi_g(k)P_g(k) - \pi_{sell}(k)P_{sell}(k))T_u(k) \tag{5.24}$$

subject to

$$0 \leq P_g(k), P_{dir}(k), P_{sell}(k) \qquad \forall k \in [1, T_b] \tag{5.25}$$

$$P_{dir}(k) + P_d(k) = P_L(k) + P_{sell}(k) \qquad \forall k \in [1, T_b] \tag{5.26}$$

$$E(0) = U \tag{5.27}$$

$$E(k) = E(k-1) + \eta_c P_c(k) T_u(k) - \frac{P_d(k)}{\eta_d} T_u(k) \qquad \forall k \in [1, T_b] \tag{5.28}$$

$$0 \leq P_c(k) + P_{dir}(k) \leq P_S(k) + P_g(k) \qquad \forall k \in [1, T_b] \tag{5.29}$$

$$0 \leq P_c(k) \leq I(k)\alpha_c \qquad \forall k \in [1, T_b] \tag{5.30}$$

$$0 \leq P_d(k) \leq (1 - I(k))\alpha_d \qquad \forall k \in [1, T_b] \tag{5.31}$$

$$I(k) \in \{0, 1\} \qquad \forall k \in [1, T_b] \tag{5.32}$$

$$u_1 P_d(k) + v_1 \leq E(k+1) \leq u_2 P_c(k) + v_2 \qquad \forall k \in [1, T_b] \tag{5.33}$$

where $I(k)$ is a Boolean integer which prevents simultaneous charging and discharging. It can be shown that the optimal solution will not have simultaneous charging and discharging due to the inefficiencies in the battery [48]. This means that the binary integer $I(k)$ may be removed, which results in a linear program (LP).

**Peak-demand pricing**

The formulation for peak-demand pricing is similar to that of time-of-use pricing. We denote $\Gamma$ to be the peak threshold, and $a(k)$ to refer to the power exceeding $\Gamma$ in time slot $k$. $\pi_{base}$ refers to the base grid price per kWh, and $\pi_{add}$ refers to the price for energy which exceeds $\Gamma$. The formulation of the control optimization problem for peak-demand pricing is as follows.

Given the time slot budget $T_b$, the predictions $[P_S(1), \ldots, P_S(T_b)]$, $[P_L(1), \ldots, P_L(T_b)]$; the battery parameters: $\alpha_c$, $\alpha_d$, $\eta_c$, $\eta_d$, $u_1$, $u_2$, $v_1$, $v_2$; the current energy content $U$; the selling price $[\pi_{sell}(1), \ldots \pi_{sell}(T_b))$, the peak threshold $\Gamma$; the prices $\pi_{base}$, $\pi_{add}$, as well as $[T_u(1), \ldots, T_u(T_b)]$:

$$\min_{\substack{P_{dir}(k),P_c(k),P_d(k),\\ P_g(k),P_{sell}(k),a(k),I(k)}} \sum_{k=1}^{T_b} (\pi_{base}(k)P_g(k) + \pi_{add}a(k) - \pi_{sell}(k)P_{sell}(k))T_u(k) \tag{5.34}$$

subject to

Constraints 5.25 - 5.33

$$a(k) \geq 0 \qquad \forall k \in [1, T_b] \tag{5.35}$$

$$a(k) \geq P_g(k) - \Gamma \qquad \forall k \in [1, T_b] \tag{5.36}$$

These optimization problems are used as part of an MPC, which performs iterative online optimization of system control over a finite time horizon which depends on the values assigned to $(T_u(k))$. In each iteration, the optimization problem is solved to determine the control actions that optimize the objective function over a finite time horizon, which is split up into $T_b$ time slots. After implementing the control action of the first time slot which corresponds to the current time, the process is repeated with an updated energy content $U = E(1)$. $P_L(1)$ and $P_S(1)$ represent the measured load and PV generation, respectively, while $[P_L(2) \ldots P_L(T_b)]$ and $[P_S(2) \ldots P_S(T_b)]$ are predictions. The forecast for each time slot is the average value expected over the corresponding time interval.

### 5.4.3 Numerical Evaluation

Two-stage and exponential timescale MPC are evaluated using load and PV generation data from ten houses in the Pecan Street Dataport [1]. The dataset consists of measurements recorded every minute over a period of one year. We use the data to simulate system

operation with mixed-timescale MPC under different configurations, varying $\zeta$ and $\delta$, and $\beta$ timescale parameters to observe the resulting cost of system operation. In the majority of our experiments, the time slot budget is fixed to 30 time slots. Mixed-timescale MPC is compared with simple MPC where the timescale consists of 30 minute-long time slots and the performance of which is indicated as a red line in the figures, as well as with the offline optimal control shown as a black line in the figures. The results of the comparison for three representative homes are analyzed in this section.

We study best-case performance by using perfect predictions for PV generation and load, deferring the use of forecasting models for future work. For consistency, the battery in each home is modelled as a 5 kWh NMC storage system whose model parameters are specified in Table 4.2. The PV generation of each home is scaled to match a 4 kW installation. Experiments with a 10 kWh battery are also shown.

**Exponential vs. two-stage timescale**

Figure 5.13 shows the resulting cost of operating the system over a period of two months with a ToU grid pricing scheme, which has a \$0.16 per kWh cost between 8 am and 8 pm, and a \$0.04 per kWh cost between 8 pm and 8 am [177]; the price of selling electricity is set to \$0.03 per kWh, which is just under the lowest grid price. Two-stage and exponential timescale MPC are both able to achieve near-optimal cost of operation. We observe that the performance of MPC with exponential timescales has low sensitivity to the specific value of $\beta$; as long as the corresponding time horizon exceeds 10-12 hours, increasing the value of $\beta$ has little effect on performance. In other words, the exponential time scale is robust to the choice of $\beta$, as long as the horizon exceeds a minimum value which in this case is 12 hours. The 12-hour mark is observed as the knee-point in the horizon vs. cost curve shown in Figure 5.10. MPC with a two-stage timescale also achieves high performance at different configurations of $\zeta$ and $\delta$, with a similar requirement that the horizon should not be shorter than 12 hours.

Figure 5.14 shows a similar comparison under peak-demand grid pricing, with a \$0.25 per kWh base price and a \$1.75 per kWh demand price for energy which exceeds the 80th percentile of that home's power demands; the price of selling electricity is set to \$0.12 per kWh. As with ToU pricing, the exponential time-scale is robust to the value of $\beta$ as long as it exceeds 6 hours, which is the knee-point in the horizon vs. cost curve shown in Figure 5.11. The two-stage timescale shows some sensitivity to the parameters, for example in Figure 5.14b where the combination of $\zeta = 1$ and $\delta = 25$ minutes has worse performance than with $\delta = 15$ minutes, despite covering a longer time horizon. This is an example of the trade-off between horizon resolution and horizon length.

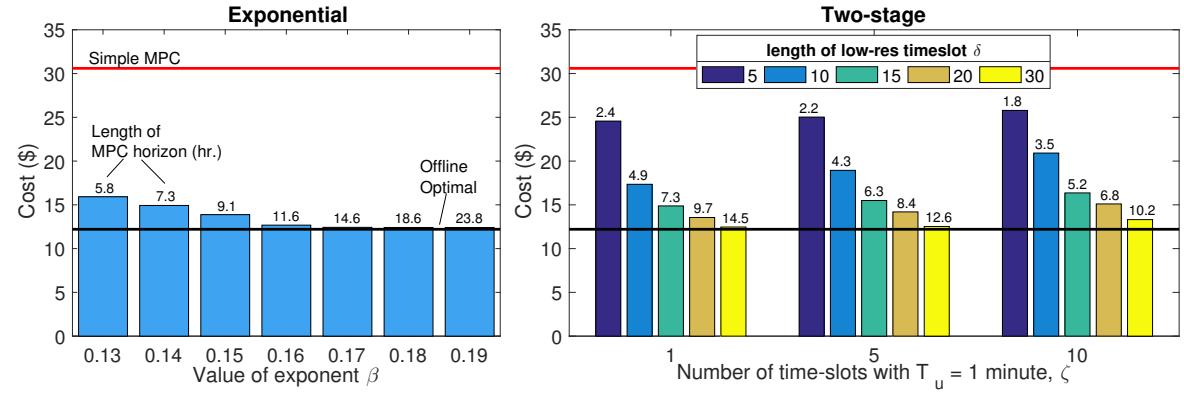


(a) House 1

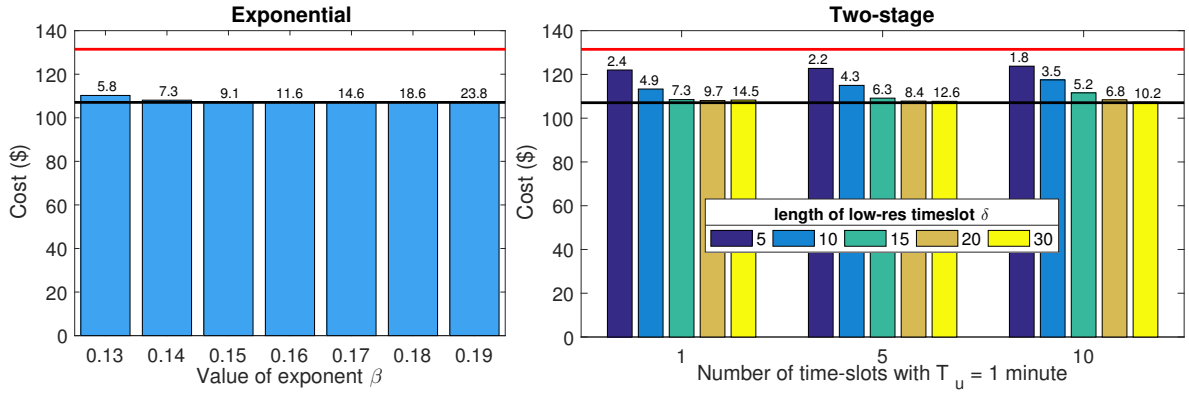


(b) House 2

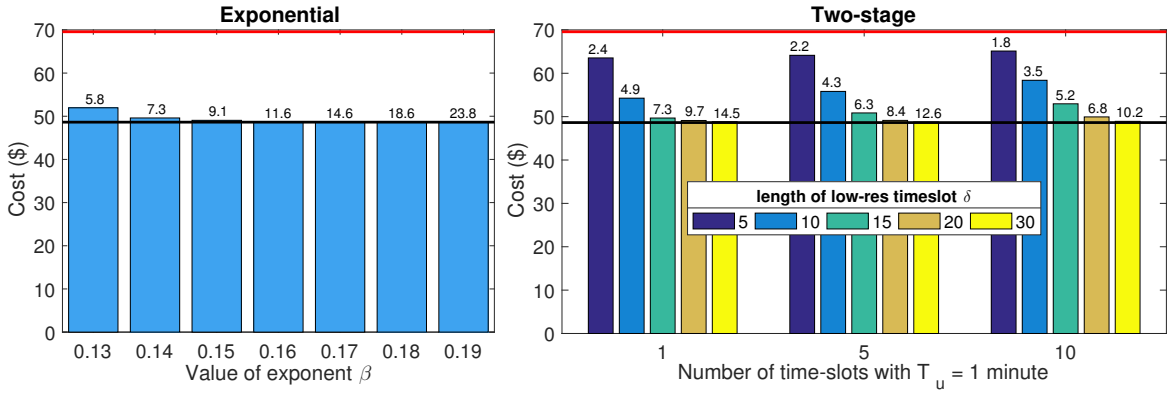


(c) House 3

Figure 5.13: Comparison of system operating cost with ToU grid pricing over a two month period using exponential, two-stage timescale MPC, basic MPC, and offline optimal.

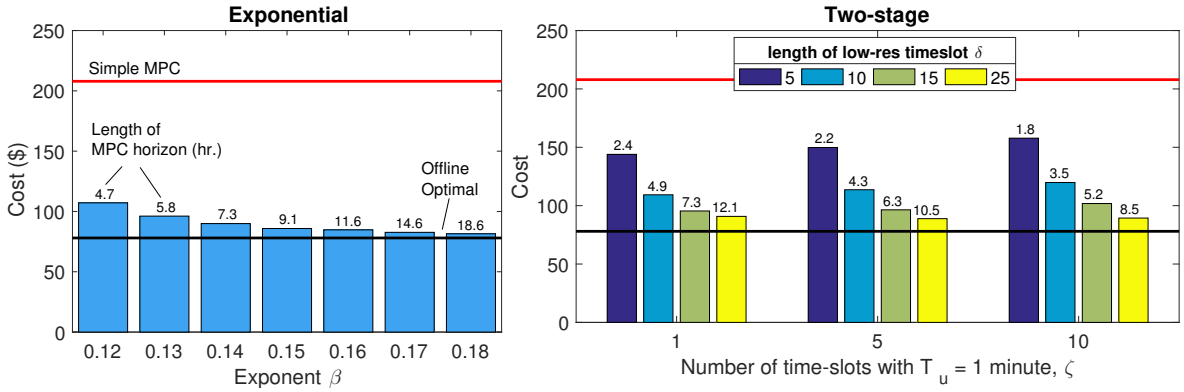


(a) House 1

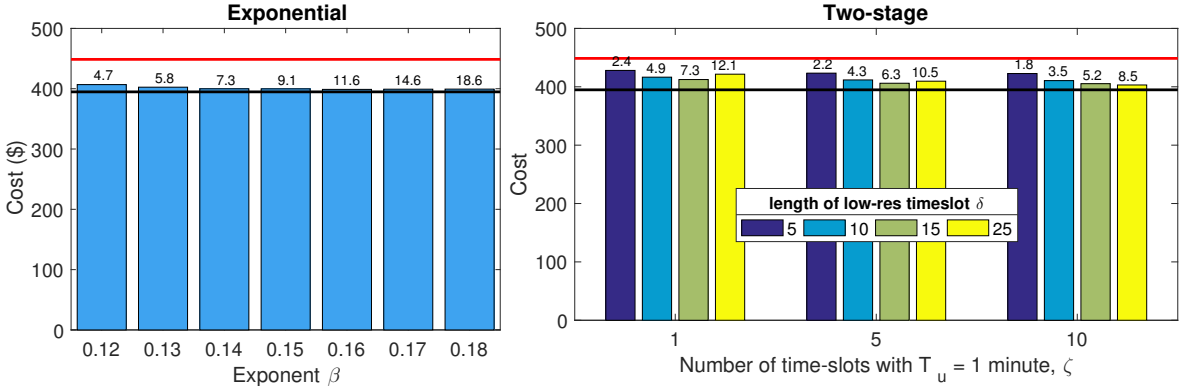


(b) House 2

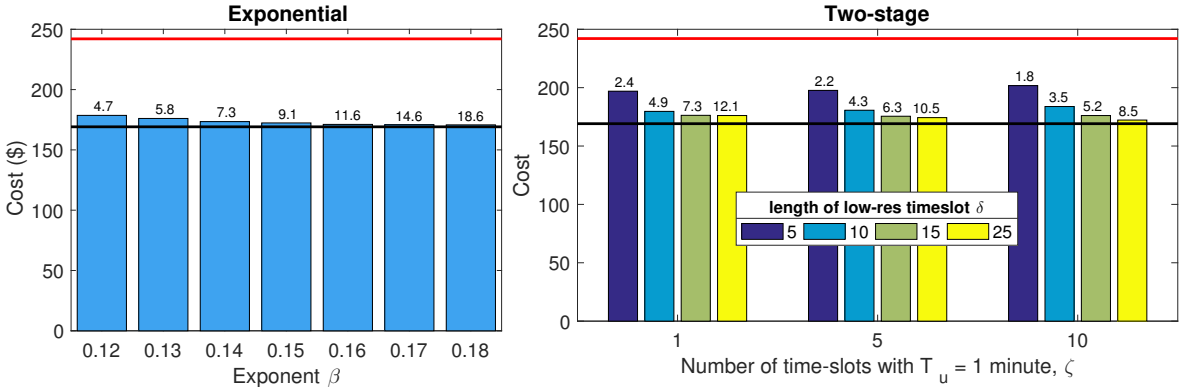


(c) House 3

Figure 5.14: Comparison of system operating cost with peak-demand grid pricing over a two month period using exponential, two-stage timescale MPC, basic MPC, and offline optimal.

Further simulations confirmed that, for a given set of parameters, the performance of MPC with both exponential and two-stage timescales remains stable over time. In other words, if a particular configuration is determined to be effective based on simulations run on data for a given month, then the same configuration is likely to be effective for the following months as well. Therefore, there is no need to reconfigure the parameters over the course of a year. The exponential timescale seems to be more robust in the selection of parameters; this observation, in combination with needing only a single configurable parameter, gives it a practical advantage over the two-stage timescale.

### Time slot budget

Using an exponential timescale, we ran simulations where the budget is restricted to 10, 20, and 30 time slots to compare the performance that may be achieved under these restrictions. A comparison of performance is shown in Figure 5.15 for ToU pricing and Figure 5.16 for peak-demand pricing. For both pricing schemes, even a 10 time slot MPC can achieve near-optimal performance with a proper configuration of $\beta$. As discussed in the preceding paragraphs, near-optimal performance is achieved when the time horizon is sufficiently long. The length of the time horizon is dependent on $\beta$ and on the time slot budget.

The results for 20 time slots further highlight the robustness of the exponential timescale. For ToU pricing (Figure 5.15), setting $\beta = 0.6$ results in a horizon which covers approximately 20 weeks, and the performance is nearly identical to $\beta = 0.28$ where the horizon is 14.1 hours. For peak-demand pricing (Figure 5.16), minor performance degradation occurs between $\beta = 0.28$ and $\beta = 0.5$, the latter corresponding to a horizon of over 3 weeks.

Figure 5.17 compares simple MPC, exponential-timescale MPC, and offline optimal grid cost with a system where the battery capacity is increased to 10 kWh. MPC results for ToU pricing are computed with a budget of 10 time slots, with $\beta = 0.7$ which gives an optimization horizon of approximately 18 hours. MPC results with peak-demand pricing are computed with a budget of 20 time slots, with $\beta = 0.32$ which gives an optimization horizon of approximately 19 hours.

### Evaluation summary

Aggregation of time slots has different consequences on the operation for ToU and peak-demand pricing schemes. For ToU pricing, the optimal strategy shifts the net energy consumption of the system to low-price periods; the amount of energy that needs to be

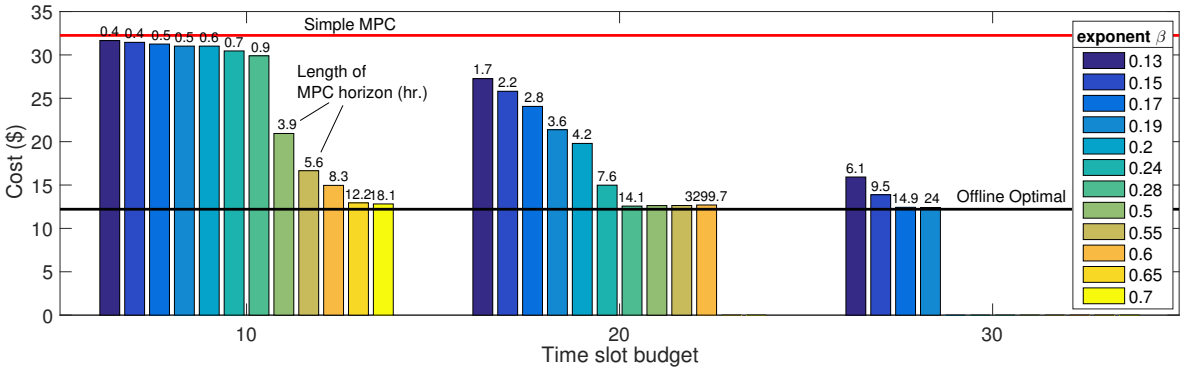


(a) House 1

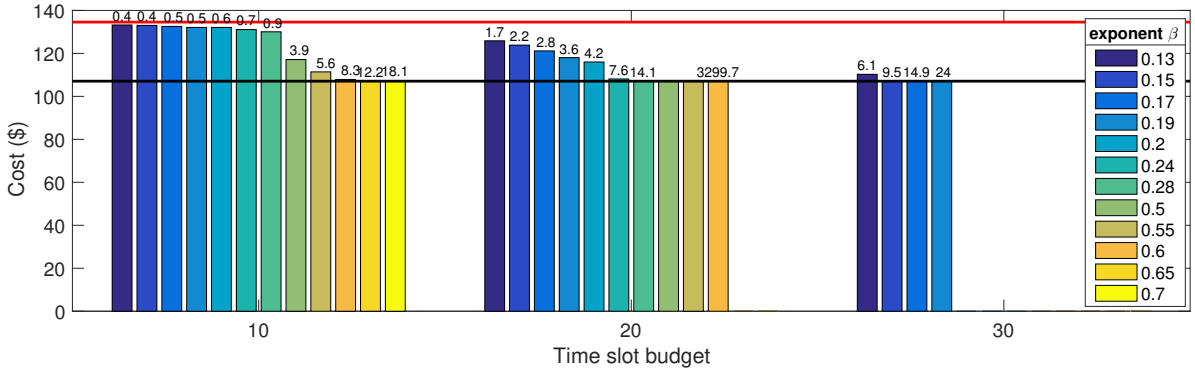


(b) House 2

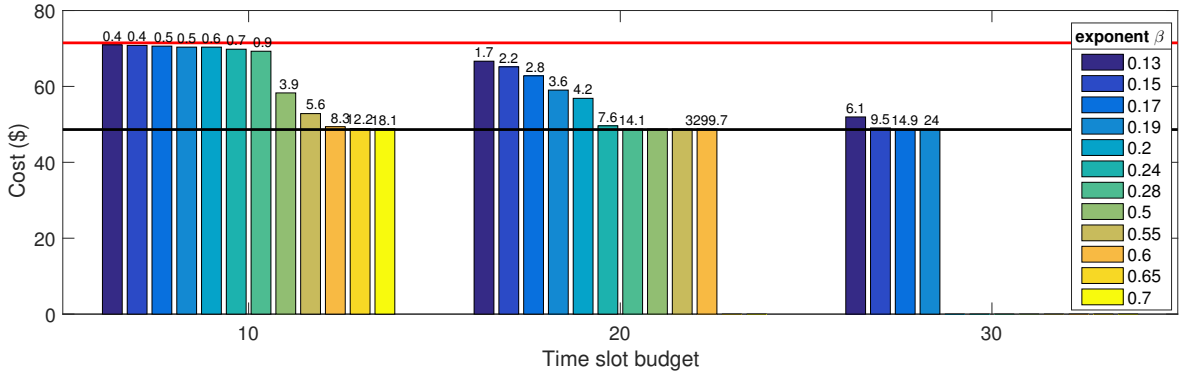


(c) House 3

Figure 5.15: Comparison of system operating cost with ToU grid pricing over a two month period and varying time slot budgets.

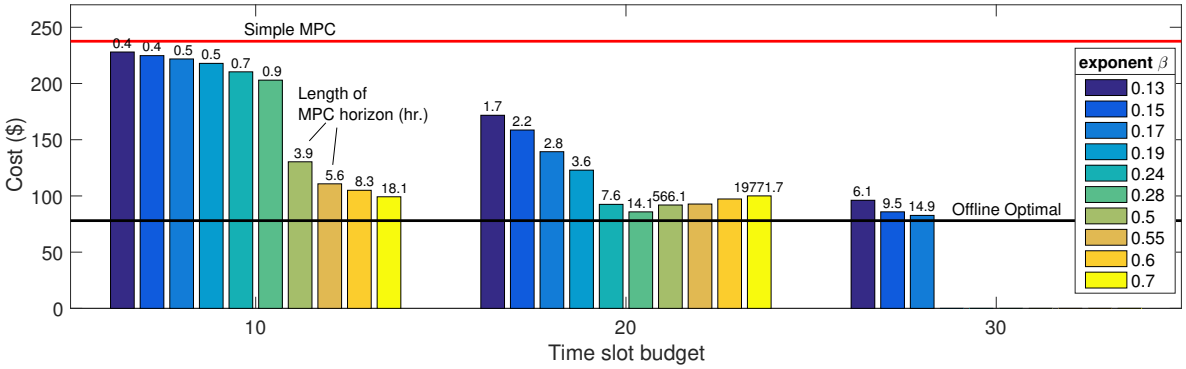


(a) House 1

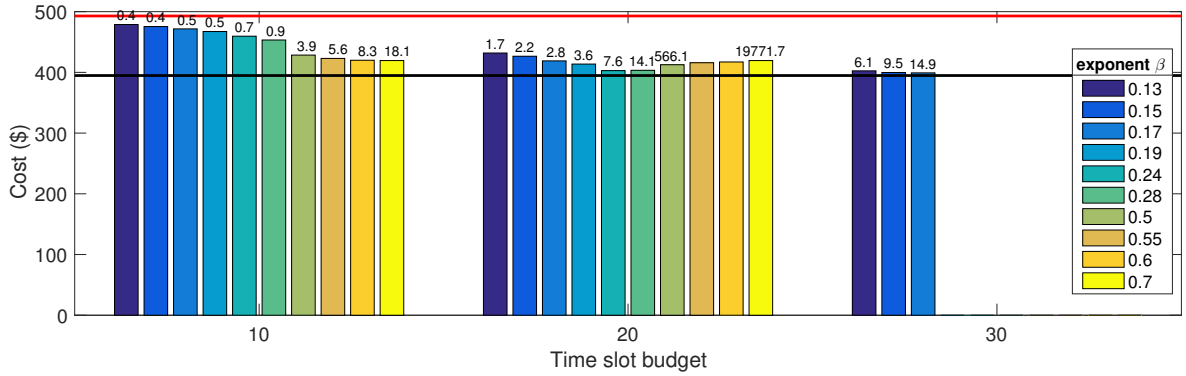


(b) House 2

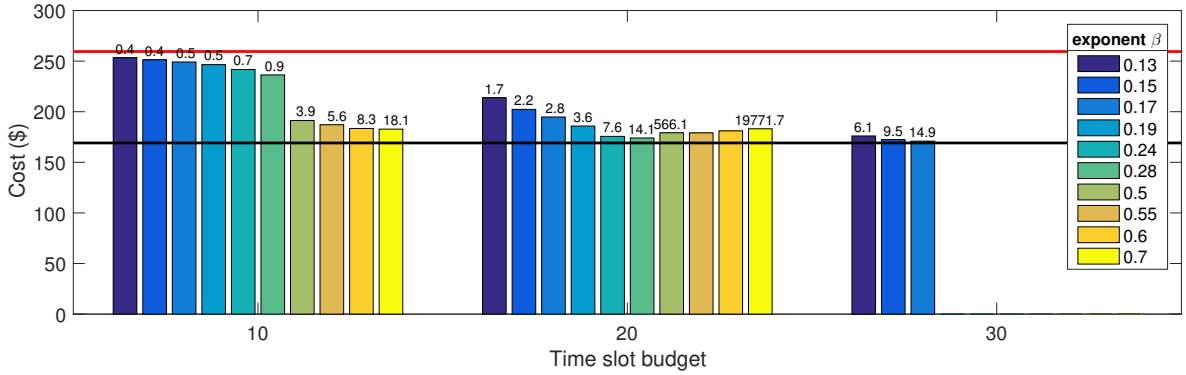


(c) House 3

Figure 5.16: Comparison of system operating cost with peak-demand grid pricing over a two month period and varying time slot budgets.

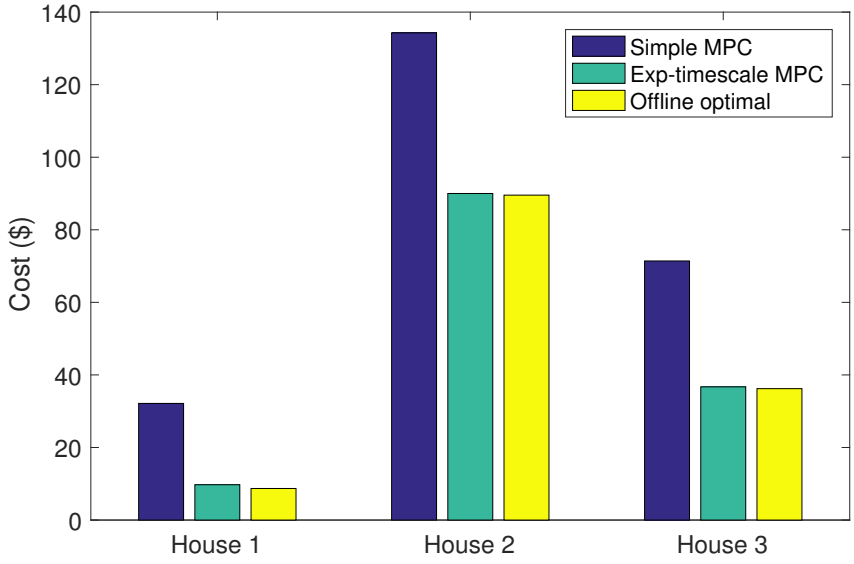


(a) ToU pricing

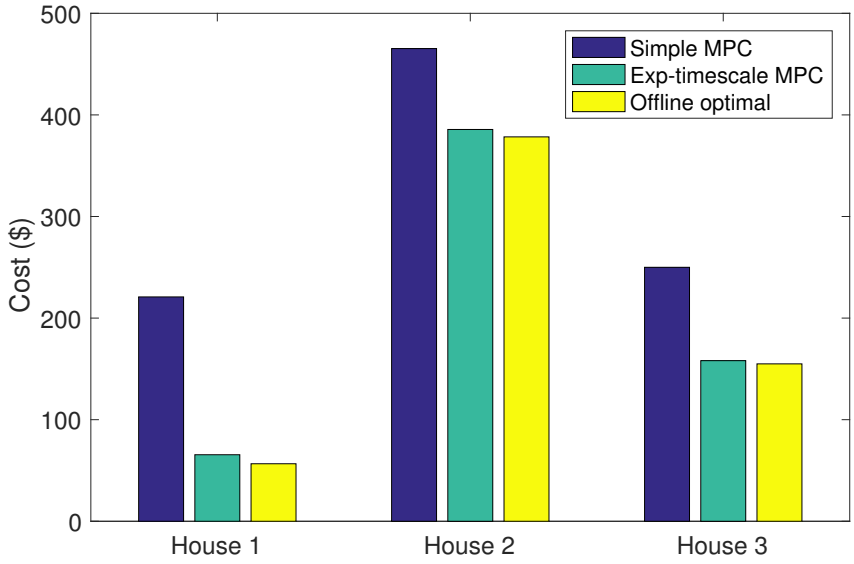


(b) Peak-demand pricing

Figure 5.17: Comparison of simple MPC, exponential-timescale MPC, and offline optimal, with a battery capacity of 10 kWh.

shifted is of key importance, and this value is not affected by the aggregation of predictions into long time slots. For peak-demand pricing, the effect of aggregating time slots is negative, because the predicted peaks in power consumption may be lost where the time horizon is aggregated into longer time slots. This effect is noticeable in our evaluation, where aggressively long horizons did not result in any performance penalties for ToU pricing, while minor penalties for peak-demand pricing were observed. However, the length of the horizon at which degradation in performance becomes noticeable is much longer than the scope of any realistic forecasting model, which practically limits the value of $\beta$ long before performance starts to degrade.

The numerical evaluation of mixed-timescale MPC leads to the following conclusions. Firstly, MPC with an exponential timescale achieves near-optimal operation of PV-storage systems under both ToU and peak-demand pricing. Secondly, configuring the parameter of the exponential timescale, $\beta$, is simple: it may take a reasonably liberal value to obtain a longer time horizon without any significant drop in operating efficacy. Thirdly, high performance with mixed-timescale MPC is possible even upon a very limited time slot budget, which implies inexpensive computational hardware.

Our evaluation was conducted with perfect forecasting models for both load in PV generation. This isolates the effects of the timescale without the added confounding effects of prediction errors from a specific practical forecasting model. However, a perfect forecasting model is not practical, and we expect the performance of MPC to decrease when imperfect predictions are used.

## 5.5 Neural Network Control

The current practice for residential battery system control is to deploy the system with a static control policy, with no further modifications over the service life of the system. However, batteries have a long lifetime, with many warranties lasting 5-10 years. The operating environment may change over the system's lifetime, which may decrease the effectiveness of the system. For example, stacking of additional applications, or changes to the grid pricing scheme[4] or to electrical load (new appliances) are possible. In such a situation, the control strategy would require adjustments to make it effective again.

Considering this dynamic operating environment, static control of the battery may be ineffective over its lifetime. Designing new rules to adapt to the changing environment is

[4]The prices and price periods in Ontario's ToU electricity pricing scheme have changed at least once a year in the recent past.

a prohibitively tedious task given the variety in the combinations of load patterns, grid pricing schemes, and applications that may require tailored control rules.

The most technologically advanced contemporary method for battery system control is MPC [159, 86, 17, 69]. An MPC may be smoothly adapted to a changing operating environment by updating its predictive models to match the changes in the patterns of load or generated power. Modifications to the grid pricing or new applications may be accommodated by updating the underlying optimization objective function and system model. However, an MPC's high computational cost may make its application to residential battery systems impractical. Although a significant decrease in computation time can be achieved by modifying the solving algorithm [178, 106], the hardware requirements of the solver for a complex non-linear control problem may still offset the effects of improvements in control for small-scale residential system applications; for such applications, a control solution with minimal online computational costs is critical. Mixed-timescale MPC, as described in Section 5.4, offers a way to achieve high performance with a tight computational budget. This section describes a complementary approach based on neural networks.

Our key insight is that, to circumvent the problem of computational cost, neural networks can be trained to approximate the control policy of an MPC and used at a fraction of the online computational cost [186, 62].

The limiting factor of using a deep neural network for residential battery system control is that training the network effectively requires data about the operating environment, which is typically unavailable at the time of system deployment. However, this data *can* be collected by the system immediately after deployment. To account for the absence of data at the time of deployment, we have designed a controller that makes periodic updates to its underlying neural network as more data is collected. We present an evaluation of our approach for a residential system comprised of PV panels and a battery, in an environment with time-of-use grid pricing, where load patterns can change over time, and with the option of selling PV power to the grid.

### 5.5.1 Design

In this section, we describe the context for our system and the controller architecture. The system configuration is described in Section 5.2 and shown in the lower half of Figure 5.18. The description of each component reflected in Figure 5.18 is provided in the subsections below.

Our design goals for the system controller are:

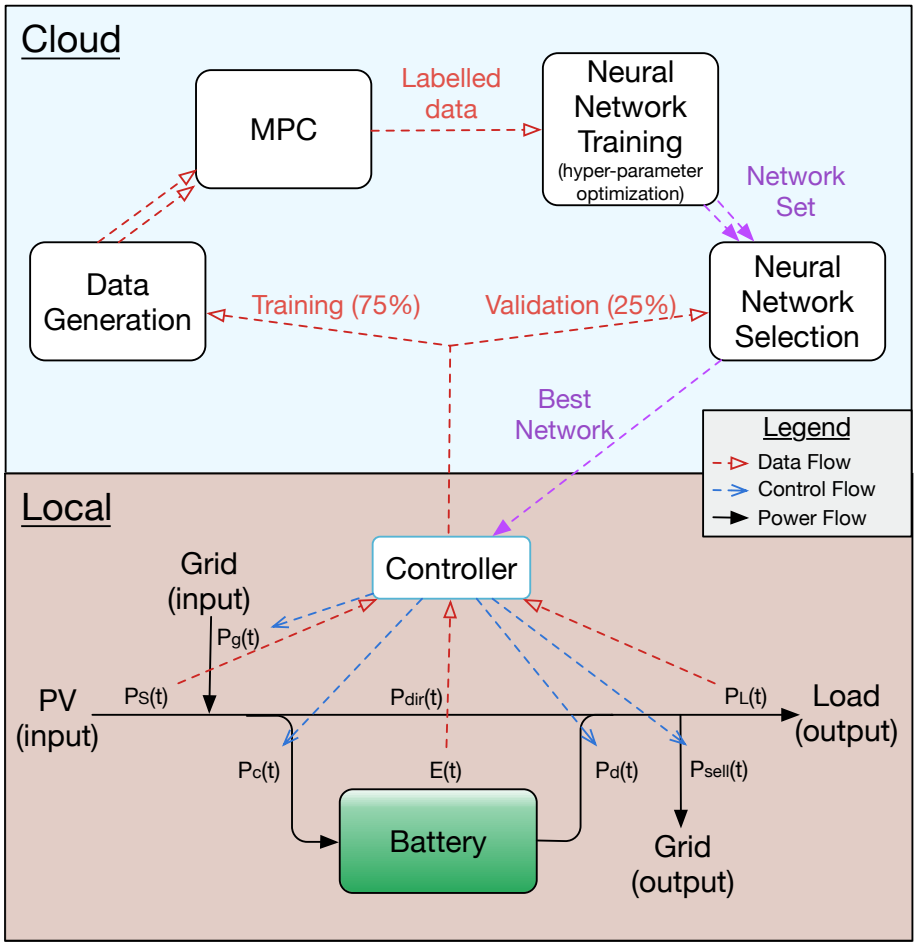


Figure 5.18: Self-learning system diagram

- **High performance**. The controller should outperform the simple control scheme (Algorithm 1, lower benchmark), and approach the performance of MPC (upper benchmark) at a fraction of the online computational cost.
- **Minimizing offline computation cost**. Due to the costs associated with cloud computing, we want to minimize the amount of offline computation.
- **Flexibility** in adaptation to different system configurations, operating environments, and applications with minimal user effort and interaction.

We note that other designs for integrating an MPC are possible, though they may be less effective. For example, one alternative would be to forego neural network approximations entirely and offload MPC computation to a cloud server; there could be a persistent connection over which local data is sent to the server, which would reply with control actions. In contrast to our proposed design, we believe this alternative approach may be costly due to server fees that would accumulate over the long lifetime of the system. In particular, our design has shown good performance in an evaluation where, after the first year of operation, communication between the local controller and the cloud server occurs only once every six months.

### Local system

The system control variables are the following:

- power used to charge the battery ($P_c$)
- power discharged from the battery ($P_d$)

Once these variables are defined, the power purchased from the grid ($P_g$) and sold to the grid ($P_{sell}$) may be directly computed to maintain the balance of power flows in the system. We assume the system has access to the grid pricing scheme, which is published by the electrical utility company.

Initially, the system is operated using a default control strategy for PV-storage systems. In our study, Algorithm 1 is used as the default strategy.

Immediately after the system is installed, it begins collecting data on solar PV generation ($P_S$) and household electricity load ($P_L$). After a set amount of time has passed, or when prompted by the system owner, the collected data is sent to a server in the cloud

and used to train a neural network. If training is successful, i.e., the network is validated to outperform the default strategy on historical observations via system simulation, then the network is uploaded into the system controller. The controller then uses the network to decide how to charge and discharge the battery, applying a post-processing step to the network output to ensure that the constraints of the system are met (see Section 5.5.1 for details). The network is retrained once a significant amount of additional data is collected. The old network is replaced if the new network is validated to outperform the old network.

### Neural network architecture

The most common type of learning studied for battery system applications is deep reinforcement learning [45, 53, 110]. However, these networks are notoriously difficult to tune, require large amounts of data and computational resources to converge to an effective solution. Our attempts with RL algorithms for this problem have been unsuccessful, often converging on a local minimum where the network prioritizes selling power to the grid rather than storing it to offset higher grid prices at a later time. Others have reported similar performance issues when studying RL control for PV-storage systems [135]. Hence, in our controller design, we forego reinforcement learning in favour of a supervised learning approach via deep feed-forward networks [54], which are trained in the cloud using an MPC to label the collected data with good control actions.

The inputs to the controller during operation are the same as those used by the MPC to compute a control action. In the case of our test system, the following inputs have been used:

- Current PV generation, $P_S(t)$;
- Current load, $P_L(t)$;
- Current battery energy content, $E(t)$;
- Grid prices[5] over time horizon $T_h$, $[\pi_g(t) \dots \pi_g(t+T_h)]$;
- PV generation predictions, $[P_S(t+1) \dots P_S(t+T_h)]$;
- Load predictions, $[P_L(t+1) \dots P_L(t+T_h)]$.

[5]For residential consumers, grid prices are typically available months in advance. We found that giving 24 hours of look-ahead was sufficient for satisfactory control performance for the system under study.

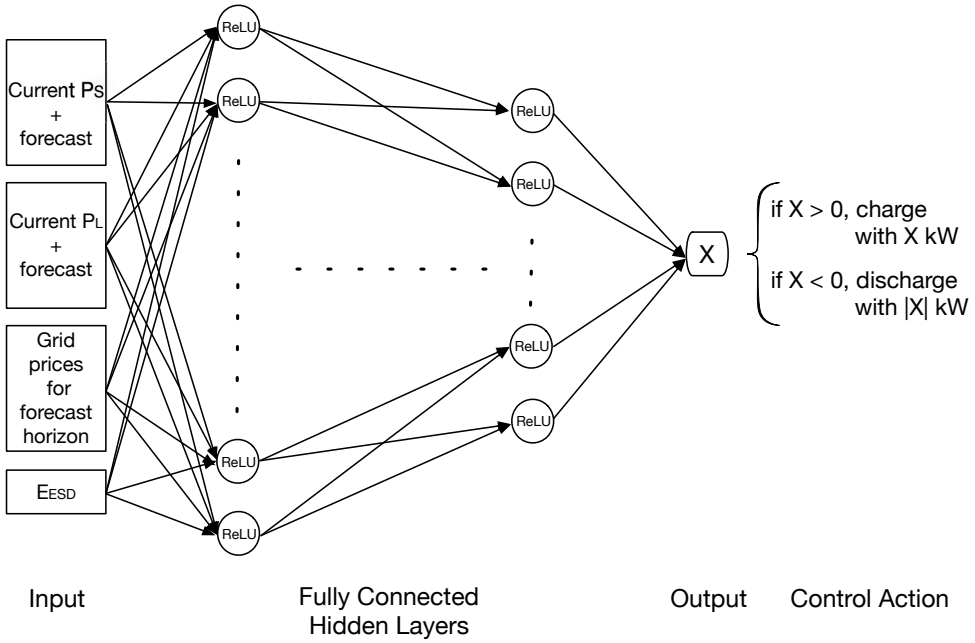


Figure 5.19: Neural network architecture

Battery control is a regression problem, where the output of the network is the charging or discharging rate (recall that once this is known, we can compute the amount of power to sell or to buy). We model the rate as a single neuron where a negative output indicates discharging and a positive output indicates charging. The magnitude of the output neuron's value is the charging/discharging rate. Figure 5.19 illustrates the network's structure.

The main hyper-parameters of the network are the number of hidden layers, the number of nodes per hidden layer, and the number of training epochs. Our training algorithm explores different combinations of hyper-parameters to find the most effective configuration. We use the rectified linear unit (ReLU) activation function for each neuron [10], and gradually decrease the number of neurons of deeper hidden layers in order to lower the training time. The output neuron does not have an activation function, which allows for the output to be a real number.

The standard practice for dealing with neural network inputs in the form of time series, such as the future grid prices and PV/load predictions, is to use convolutional filter layers to extract features from each series [128]. The cost of using 1-D convolutional layers would have to be balanced with the cost of searching through the space of other network hyper-parameters. The results presented in this section use networks which do not have

any convolutional layers, since we found that they were not necessary to achieve effective control performance for the system under study, though they may become necessary for other systems.

### Training data

Deep neural networks perform better when a large amount of training data is available. The amount of time needed to collect a sufficiently large amount of data after system installation can be prohibitive to the adaptive neural network approach. As a possible workaround, additional data may be synthesized by using model-based approaches such as ARMA, Gaussian Mixture Models, or more detailed models such as [107]. In our research, we use the following simple method to expand the amount of PV generation and load data available for training:

1. Start with original training set of size $N$;
2. Generate a vector of $N$ *perturbation factors* by conducting a bounded random walk in the range 0.95 and 1.05 with a step size of 0.01;
3. Take an element-wise product of the training set and the perturbation factors vector, and add the resulting dataset to the expanded training set;
4. Repeatedly perturb the original training set until sufficient[6] amount of training data has been generated.

The perturbed data are heavily correlated to the original dataset and do not add completely unseen patterns to the training set. Nevertheless, our approach creates altered sets of PV and load data that resemble the original measurements (within 5%) and preserves the continuous structure of data over time, while being sufficiently different to provide the network with "fresh" data to train on. In our testing, we have observed that increasing the size of the training dataset by 3-5× with this data generation approach improves the control performance of the network in the first year of system deployment where the original training dataset is small (<100 days).

Note that we do not use all of the collected data for training; 25% of the data is reserved for network selection, in which we validate the performance of the selected network. This data is labelled with good control actions using an MPC, as discussed in the next subsection.

[6]The ideal amount of generated training data may be determined by continuing to add data to the training set while validation performance is increasing.

### MPC

One of the essential aspects of our approach is labelling the training data with *effective* control actions. In general, a neural network could be used to approximate any practical control strategy. By practical, we mean that the strategy uses only attainable data to compute the control actions. The inputs of the strategy map to the inputs of the neural network. We note that using the true optimal control for labelling would be impractical because it may depend on perfect knowledge of the entire operating horizon, which is typically impossible to obtain in an online control setting.

To label the data with control actions, we formulate an MPC, i.e., a discrete mathematical optimization problem for optimal system control over a given time horizon. The details of the formulation[7] are provided in Section 5.4.2. The inputs to the neural network are the same as the inputs to the MPC. We solve the MPC repeatedly to label every time slot of solar, load, and initial battery state data with the charging and discharging actions computed by the MPC. The inputs at time slot $t$ are assigned the label $P_c(t) - P_d(t)$ as computed by the MPC. In each iteration of MPC, the battery energy content is updated to reflect the action taken in the previous iteration. The labelled dataset is used to train neural networks.

A neural network controller trained using MPC can adapt to system changes that are expressed in the PV generation and load data without any modifications to the controller. Changes to the system, grid pricing scheme, and applications are communicated to the controller via modification of the underlying optimization problem of the MPC. The MPC labels the data with updated control actions by which the neural network learns how to operate in the new environment. We note that a linear program can be solved even with basic computational hardware, and a linear MPC could be used to operate the system directly without a neural-network approximation. We view the linear MPC in our exploration of the design space as a first step in evaluating the potential of this approach, since it greatly simplifies and accelerates the analysis.

### Training

In this step, we train many different neural networks, each with a different structure, i.e., different numbers of layers and nodes per layer. Each one of the network configurations is trained using the back-propagation training algorithm with a mean-squared error cost

[7] This work was completed before our work with mixed-timescale MPC, and hence here we use a simple MPC where all time slots have equal length.

function for minimizing the difference between the network output and the MPC actions computed for the given input. During training, we use a dropout rate[8] of 0.2 at each hidden layer to avoid overfitting [152].

The network is trained with a set number of epochs (i.e., passes through the training data), and the network weights at the end of each epoch are saved. After training is completed, the set of networks is tested to select the network which will be uploaded to the controller.

### Selection

To assess the performance of the trained networks, we run a set of simulations in which each network is used to make control decisions for the system. The system model in the simulation is the same as the one used in the LP formulation described in Section 5.5.1.

The selection process consists of two steps.

Step 1: All of the newly trained networks are simulated on 25% of the data which has been with-held from training, to single out the network with the most effective performance on unseen data.

Step 2: The best network from Step 1 is compared against the network that is currently deployed, by simulating them both on all of the data collected so far; if there is no network currently deployed, the new network is compared against the default control strategy. If the new network is the winner, it is uploaded to the controller.

Note that in the simulations as well as in the controller, the output of the network is processed by an algorithm to ensure that the physical constraints of the system are met before the control action is taken. Firstly, the algorithm ensures that the constraints of the battery are being met, as described by Constraints 5.30, 5.31, and 5.33 in the LP, by decreasing the charging/discharging action if necessary. Secondly, the power in the system is balanced (Constraints 5.25 and 5.26) by selling less and then buying more power from the grid if there is a power deficit, or buying less and then selling more if there is a power surplus.

Table 5.2: Battery model parameters, per household

| Parameter | $\alpha_c$ | $\alpha_d$ | $u_1$ | $u_2$ | $v_1$ | $v_2$ | $\eta_c$ | $\eta_d^*$ |
|---|---|---|---|---|---|---|---|---|
| **House 1** | 8 | 8 | 0.053 | -0.125 | 0 | 8 | 0.99 | 1.11 |
| **House 2** | 8 | 8 | 0.053 | -0.125 | 0 | 8 | 0.99 | 1.11 |
| **House 3** | 14 | 14 | 0.053 | -0.125 | 0 | 14 | 0.99 | 1.11 |

*includes inverter efficiency penalty of ∼10%

### 5.5.2 Evaluation

We use four years of PV generation and load data obtained from three houses in the Pecan Street Dataport [1] to simulate different realizations of the system operating environment. The hourly data have been collected from Texas, USA, over a period of four years. All three homes have high electricity consumption in the range of 35-45 kWh per day, which makes them good candidates for a PV panel and battery installation. Reflecting a ToU pricing scheme offered by a Texas utility company [177], the grid price in the night time between 21:00 and 08:59 is set to \$0.04/kWh, while the price in the day time between 09:00 and 20:59 is set to \$0.16/kWh between October and May, and \$0.21/kWh between June and September. The price of selling electricity is set to a constant value of \$0.03/kWh, which is just below the lowest grid price.

We compare our adaptive controller, the default control algorithm, and the MPC in terms of the grid cost paid by the homeowner. We also compare these costs with the lowest possible (optimal) cost obtained via an oracle, which we compute by solving the LP in Section 5.5.1 over the entire four-year horizon ($T_h = 4$ years). We use perfect hourly forecasts ($T_u = 1$) with $T_h = 24$ as input for the MPC and neural network simulations, deferring an evaluation with realistic prediction errors and optimization of the prediction horizon to future work. The parameters used in our system model are given in Table 5.2. The neural networks are implemented in Python using the Keras/Tensorflow libraries and trained with the 'adam' optimizer.

To model the behaviour of a realistic system that periodically re-learns the best operating strategy by reviewing its immediate prior past, the network training process was activated on the following days: 10, 20, 40, 80, 120, 190, 365, 550, 730, 910, 1090, 1270, and 1450. The intervals between training times start off being short to take advantage of the relatively rapid growth of the training set and increase later as the novelty and importance of the new data decreases. During the data generation phase, the amount of training data

[8]Dropout rate refers to the percentage of neurons that are chosen at random to be ignored during a single iteration of the training algorithm.

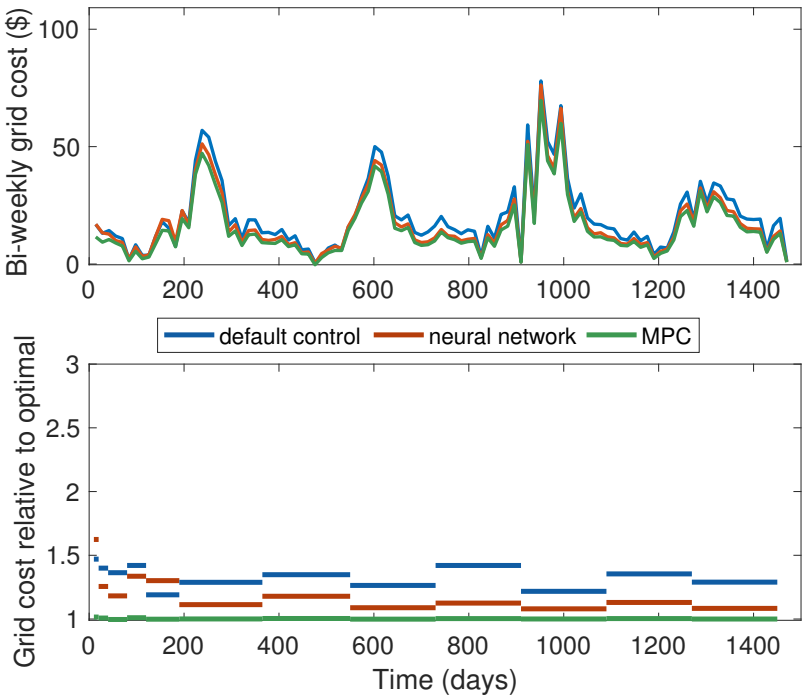


Figure 5.20: House 1: bi-weekly grid cost, and the total cost between re-training periods relative to the optimal cost.

is increased by a factor of 4 by perturbing the traces. During the training phase, we train networks with 3 to 7 hidden layers, 40 to 100 nodes per layer, and 21 epochs[9]. A total of 24 different network structures are tested at each retraining period, each with 21 different sets of weights, one for each training epoch; this results in a total of 504 different networks to choose from in the selection step. The largest network has a total of approximately 32K parameters.

The results for the three houses are summarized in Figures 5.20, 5.21, and 5.22, respectively. Each figure shows the bi-weekly grid cost over four years for the three control methods, as well as a comparison of their relative performance with respect to the oracle. Across all three houses, we see that MPC is quasi-optimal. The neural network can get within 25% of the optimal cost 200 days after deployment, and within 5-10% in 2-4 years. For houses 1 and 2, the neural network outperforms the default strategy within 100 days

[9]We determined these ranges to be sufficient for obtaining a performant network through experiment.

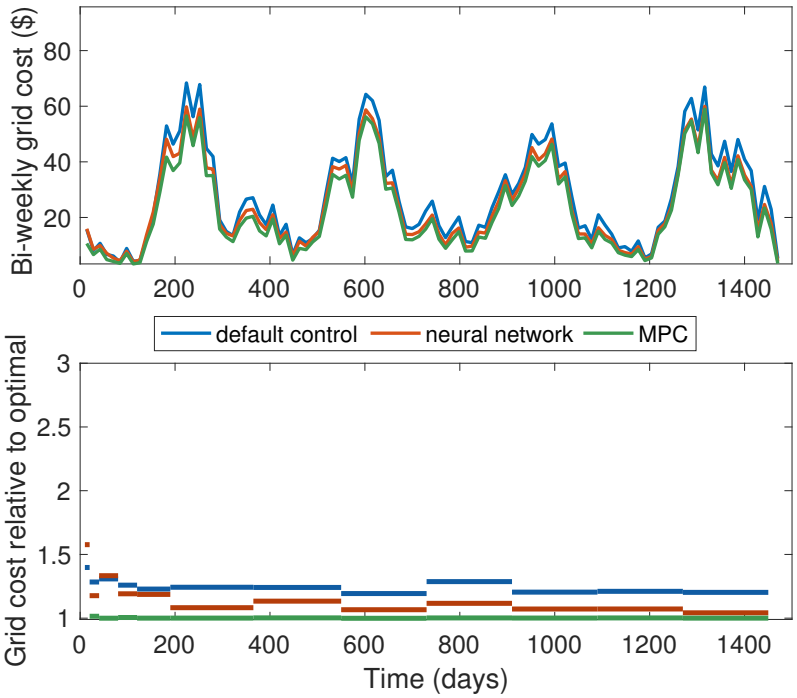


Figure 5.21: House 2: bi-weekly grid cost, and the total cost between re-training periods relative to the optimal cost.

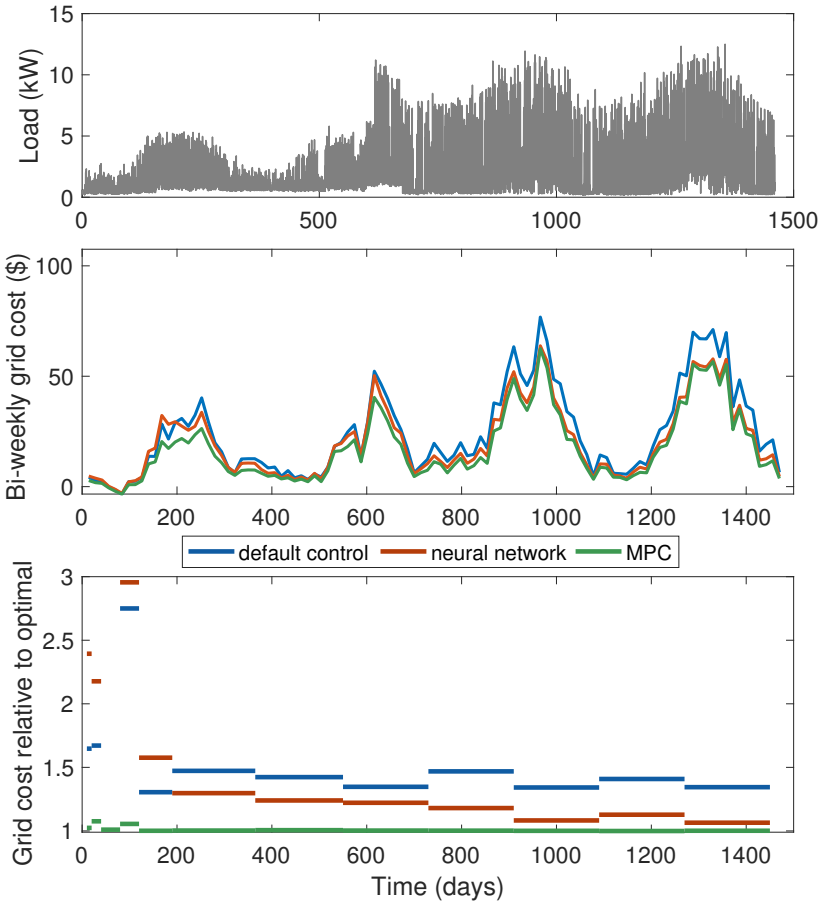


Figure 5.22: House 3: load, which has large changes halfway through year 2, the biweekly grid cost, and the total cost between re-training periods relative to the optimal cost.

of system deployment. Across all houses, the adaptive controller saved between 13-19% on grid costs over the 4 years compared to the default strategy, with savings as high as 25% for the third house after two years of data was collected.

Figure 5.22 also shows the changes in the load pattern approximately halfway through the second year of deployment. We have observed that the networks trained on data prior to the load change performed poorly after the change. This observation highlights the value of re-training the controller on recently-collected data.

We observe a general trend in the size of the most effective networks across different training periods. Smaller networks, i.e, those with fewer nodes and hidden layers, have better performance with smaller sets of available data, while larger networks eventually perform better as more training data becomes available. To reduce the amount of computation done during the training of the network set, one possible optimization is to check for convergence of the structure of the best network and remove structures with historically poor performance from the set.

In terms of online computation cost, the neural networks implemented using Python's Keras deep learning library and run on a single 1.6 GHz CPU have been on average 22× faster than a linear MPC run using the CPLEX LP solver [33] on a machine with up to 24 3.0 GHz CPU cores for parallelization and 500 GB of RAM. Using a non-linear MPC – in the case where a more accurate non-linear battery model is needed or if the grid pricing scheme or application constraints could not be expressed linearly – requires orders of magnitude more time, and may be too slow for online control using the limited hardware that is typically available in a battery controller of a residential home system.

### 5.5.3 Summary

We have developed an adaptive controller framework for PV-storage systems based on neural networks, MPC, and system simulation. Unlike the static algorithms typically used in currently-deployed battery systems, this controller may adapt to changes in the system. An evaluation of the controller with time-of-use pricing has shown that it approaches MPC-level performance at a fraction of the online computational cost of MPC.

## 5.6 Future Work on Operating Methods

In this chapter, we have developed approaches to PV-storage system operation based on rule-based algorithms, MPC, and neural networks. We reflect on the relative online com-

Table 5.3: Relative Comparison of Operation Methods

| | **MPC** | **Rule-based** | **Neural Network** |
|---|---|---|---|
| Online complexity | high | low | low |
| Offline complexity | low | high | high |
| Interpretability | med | high | low |

putational cost, offline complexity, and interpretability of each approach in Table 5.3. Each of these methods may be developed further through additional research and improvement, as discussed next.

### 5.6.1 Evaluation with Forecasting Models

In the evaluation of our control approaches, with one exception[10], we have assumed perfect predictions of load and PV generation processes over a finite time horizon. Disregarding prediction errors allowed us to decouple the theoretical performance of a control method from the accuracy of forecasting models. Nevertheless, taking into account prediction errors is important in practice. There are no known models that can perfectly forecast load and PV generation, and errors in the forecast can be expected to affect the performance of control methods that rely on accurate predictions [70]. Accounting for forecasting errors through the use of robust optimization approaches such as stochastic MPC [126] is a potential direction for future work. However, using stochastic optimization approaches may increase the computational complexity over regular optimization by an order of magnitude [187], which serves as further motivation for our work on mixed-timescale MPC and neural network approximations.

### 5.6.2 Verification on Complex Systems

The control methods described in this chapter have been evaluated on systems that reflect real operating deployments. In the development of our methods, which are based primarily on optimization models, we have paid specific attention to the online computational complexity. This has been done to allow our methods to be used even in cases where the optimization problem is non-linear. Nevertheless, our evaluation has been restricted to systems that could be adequately represented via linear models, which has allowed us to

[10]The evaluation of the day-ahead X strategy for time-of-use pricing was tested under the assumption of a Gaussian error

conduct a wide range of case studies within reasonable time frames. Extending the evaluation to non-linear systems, such as those where a more accurate and non-linear battery model is necessary or where the objective function is non-linear, would further validate our control methods and may lead to further research and development of approaches addressing these challenges in system complexity.

**Chapter 5 Summary**

We study the problem of optimizing PV-storage system operation, with the objective of minimizing the cost of meeting the system owner's load. We develop control strategies for PV-storage systems under two grid pricing schemes: time-of-use pricing and peak-demand pricing. Among our contributions is (a) an algorithm for each pricing scheme which achieves high performance without requiring any forecasting, (b) a mixed-timescale model predictive control approach which achieves quasi-optimal performance with relatively low computational hardware requirements, and (c) a deep neural network controller which is capable of learning an effective control policy over time and adapting to changes in the operating environment.

# Chapter 6

# Summary and Conclusion

PV-storage systems offset the grid energy consumption of end-consumers with a clean energy source and add flexibility to their consumption process. The high capital cost of these systems has motivated the study of how to increase their value to facilitate their adoption. The focus of this work has been the increase of PV-storage system value through the optimization of system design and operation. In this chapter, we summarize our specific contributions in Section 6.1, discuss the integration of our work in the solution of related problems in Section 6.2, and conclude the thesis in Section 6.3.

## 6.1 Summary of Contributions

Physical system models, together with data traces that capture the dynamic operating environment, have been at the center of our approach to system optimization. Each of these components has been analyzed and developed in this dissertation.

We propose the PI model that allows users to accurately simulate the dynamic behaviour of Li-ion batteries. The PI model has several distinctive features. Firstly, its calibration is based on a battery spec sheet without the necessity of any complex experiments. Secondly, it integrates both the cell and BMS components into a single model. Thirdly, it uses power, as opposed to current, as model input. These features significantly simplify the task of running energy system simulations without compromising on model accuracy, as demonstrated through an extensive set of experiments. It has been publicly released for use on Matlab's Simulink platform.

We derive the C/L/C, C/L/L, and L/L/Q models out of the PI model. They are analytic and represent different combinations in the trade-off between model accuracy and complexity. This allows the researcher to select the appropriate model depending on the objective and the operating range of the given battery optimization study, which increases the power of the research through a better correspondence between the model and its application.

Using a battery model in conjunction with a set of data traces, we have developed a method for PV-storage system sizing which takes into account the uncertainty in the degree to which the available historical data reflects the future operating environment. This method uses system simulations and an empirical probability bound to provide sizing recommendations that meet the objectives of the system owner at relatively low cost.

Specialized algorithms for system operation developed in this work are based on heuristics obtained through analysis of offline optimal control over a set of data traces and reduce the operating cost of PV-storage systems. They achieve significant cost savings for time-of-use pricing and peak-demand pricing schemes. In particular, for each pricing scheme, a myopic algorithm has been developed. This substantially reduces the implementation complexity while also achieving low operating costs.

The model predictive control algorithms with heterogeneous timescales developed in this research are capable of achieving quasi-optimal operating costs for PV-storage systems with limited computational resources. It has been empirically shown that an exponentially growing timescale achieves high performance for ToU and peak-demand pricing while requiring the calibration of only one parameter.

An adaptive neural network controller has been developed for ToU pricing. It is capable of learning an efficient operating strategy over time and reacting to changes in the operating environment. This controller is useful when the control problem is too complex for online optimization with MPC.

## 6.2 Future Work

In the final paragraphs of Chapters 3, 4, and 5, we have discussed specific directions for the further development of each topic. Our models can be improved by taking into account temperature and degradation effects; system sizing may be studied with a stricter assumption on the availability of data; our operating methods could be integrated with forecasting models and evaluated on complex non-linear systems. Here, we briefly discuss how our contributions may be integrated as a part of a wider study of energy systems.

The battery models developed in this dissertation are not restricted to the analysis of PV-storage systems. They may be used in other applications as well, including electric vehicles and drones. A large-scale optimization study involving thousands of electric vehicles would have many optimization variables, and would thus require models with low complexity, such as the C/L/C model, for the problem to be solvable in a reasonable amount of time. A similar case can be made of the optimization of battery-powered autonomous aircraft, where the problems of coordinated charging and flight routes require a simple and accurate battery model.

In a large-scale multi-system analysis of PV-storage systems, the control strategies developed in this thesis may be used to study the combined effects of multiple systems on the operation of the grid. It is possible that many systems individually optimizing their operating costs may collectively have a negative effect on grid stability. The resolution of this potential problem may require the development of grid pricing schemes to properly motivate individual systems to have grid-friendly behaviour. Our exponential timescale MPC allows for the exploration of this problem at scale.

## 6.3 Remarks

To conclude this thesis, we briefly discuss how we address some of the research challenges outlined in Section 1.1. Specifically, we elaborate on the challenges of limited data and system complexity.

In today's world, the typical challenge in data-driven studies is to process large quantities of data and derive some general or specific conclusions. However, in the data-driven study of PV-storage systems, which are not yet widely adopted, the challenge is quite the opposite: how to extract value from the limited available data. This was one of the major aspects of the dissertation. The developed battery models are calibrated entirely on easily-obtainable data; our robust sizing framework addresses insufficient data directly through a statistical approach; our neural network controller learns as data becomes available.

Optimization of complex non-linear systems is a challenging task requiring lots of computational power. Many of the practical limitations of contemporary research problems are caused by a lack of computational power, including the problems addressed in this thesis. Our battery models are designed to balance accuracy and computational complexity; our work on robust sizing includes a direct comparison of multiple approaches in terms of their computations time; our work on MPC timescales and neural network approximations is a direct reflection of the limited computational capabilities of embedded controllers. However, the progress of computational resources is an obvious and continuing trend. It will

unavoidably lead to a relaxation of these limitations, and correspondingly increase the significance of the power of computation to support energy system applications.

On one hand, the increase of computational power will simplify algorithms, because the reason for making every last micro-optimization to conserve processor cycles weakens. On the other hand, this increase will open new doors for the digitization and optimization of systems which are not currently feasible, but may be feasible tomorrow. Accordingly, the need for reliable, performant algorithms is ever-growing. This dissertation is an instance of how systems that have not been digitized in the past become digitized in the present.

# References


[1] Pecan Street Inc. Dataport, 2018.

[2] A123 Systems. *High Power Lithium Ion APR18650M1A*, 2009. LiFePO4 cell specifications.

[3] Sherif Abdelrazek and Sukumar Kamalasadan. A novel integrated optimal battery energy management control architecture considering multiple storage functions. In *North American Power Symposium (NAPS), 2014*, pages 1–6. IEEE, 2014.

[4] Adedamola Adepetu and Srinivasan Keshav. Understanding solar pv and battery adoption in ontario: An agent-based approach. In *Proceedings of the Seventh International Conference on Future Energy Systems*, e-Energy '16, pages 5:1–5:12, New York, NY, USA, 2016. ACM.

[5] US Energy Information Administration. Us electricity generation. https://www.eia.gov/tools/faqs/faq.php?id=427&t=3, 2018.

[6] US Energy Information Administration. Us energy storage market trends. https://www.eia.gov/analysis/studies/electricity/batterystorage/pdf/battery_storage.pdf, 2018.

[7] Kodjo Agbossou, Mohanlal Kolhe, Jean Hamelin, and Tapan K Bose. Performance of a stand-alone renewable energy system based on energy storage as hydrogen. *IEEE Transactions on Energy Conversion*, 19(3):633–640, 2004.

[8] Andreas Aichhorn, Michael Greenleaf, H Li, and J Zheng. A cost effective battery sizing strategy based on a detailed battery lifetime model and an economic energy management strategy. In *Power and Energy Society General Meeting, 2012 IEEE*, pages 1–8. IEEE, 2012.

[9] Omid Ardakanian, Catherine Rosenberg, and S. Keshav. On the impact of storage in residential power distribution systems. *SIGMETRICS Perform. Eval. Rev.*, 40(3):43–47, January 2012.

[10] Raman Arora, Amitabh Basu, Poorya Mianjy, and Anirbit Mukherjee. Understanding deep neural networks with rectified linear units. *arXiv preprint arXiv:1611.01491*, 2016.

[11] National Hydropower Association. Comments of the national hydropower association on the miso psc discussion of sata modeling and analysis. https://www.hydro.org/wp-content/uploads/2019/01/NHA-MISO-Comments-1.7.2019.pdf, 2019.

[12] Raji Atia and Noboru Yamada. Sizing and analysis of renewable energy and battery systems in residential microgrids. *IEEE Transactions on Smart Grid*, 7(3):1204–1213, 2016.

[13] P Attaviriyanupap, K Tokuhara, N Itaya, M Marmiroli, Y Tsukamoto, and Y Kojima. Estimation of photovoltaic power generation output based on solar irradiation and frequency classification. In *2011 IEEE PES Innovative Smart Grid Technologies*, pages 1–7. IEEE, 2011.

[14] B. Azzopardi and J. Mutale. Smart integration of future grid-connected pv systems. In *34th IEEE Photovoltaic Specialists Conference (PVSC)*, pages 2364 – 2369, June 2009.

[15] Anirudh Badam, Ranveer Chandra, Jon Dutra, Anthony Ferrese, Steve Hodges, Pan Hu, Julia Meinershagen, Thomas Moscibroda, Bodhi Priyantha, and Evangelia Skiani. Software defined batteries. In *Proceedings of the 25th Symposium on Operating Systems Principles*, pages 215–229. ACM, 2015.

[16] S Bahramara, M Parsa Moghaddam, and MR Haghifam. Optimal planning of hybrid renewable energy systems using homer: A review. *Renewable and Sustainable Energy Reviews*, 62:609–620, 2016.

[17] Kyri Baker, Junyao Guo, Gabriela Hug, and Xin Li. Distributed mpc for efficient coordination of storage and renewable energy sources across control areas. *IEEE Transactions on Smart Grid*, 7(2):992–1001, 2016.

[18] Raul Banos, Francisco Manzano-Agugliaro, FG Montoya, Consolacion Gil, Alfredo Alcayde, and Julio Gómez. Optimization methods applied to renewable and sustainable energy: A review. *Renewable and sustainable energy reviews*, 15(4):1753–1766, 2011.

[19] L. Barra, S. Catalanotti, F. Fontana, and F. Lavorante. An analytical method to determine the optimal size of a photovoltaic plant. *Solar Energy*, 33(6):509 – 514, Sep 1984.

[20] Robert Basmadjian, Yashar Ghiassi-Farrokhfal, and Arun Arvishwa Keshav. Hidden storage in data centers: Gaining flexibility through cooling systems. In *19th International GI/ITG Conference on Measurement, Modelling and Evaluation of Computing Systems*, 2018.

[21] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.

[22] Arnab Bhattacharya, Jeffrey Kharoufeh, and Bo Zeng. Managing energy storage in microgrids: A multistage stochastic programming approach. *IEEE Transactions on Smart Grid*, 2016.

[23] Ted KA Brekken, Alex Yokochi, Annette Von Jouanne, Zuan Z Yen, Hannes Max Hapke, and Douglas A Halamay. Optimal energy storage sizing and control for wind power applications. *IEEE Transactions on Sustainable Energy*, 2(1):69–77, 2011.

[24] Daniel Burmester, Ramesh Rayudu, Winston Seah, and Daniel Akinyele. A review of nanogrid topologies and technologies. *Renewable and Sustainable Energy Reviews*, 67:760–775, 2017.

[25] Giacomo Capizzi, Francesco Bonanno, and Giuseppe M Tina. Recurrent neural network-based modeling and simulation of lead-acid batteries charge–discharge. *IEEE Transactions on Energy Conversion*, 26(2):435–443, 2011.

[26] Manuel Castillo-Cagigal, Estefanía Caamaño-Martín, Eduardo Matallanas, Daniel Masa-Bote, Álvaro Gutiérrez, Felix Monasterio-Huelin, and Javier Jiménez-Leube. Pv self-consumption optimization with storage and active dsm for the residential sector. *Solar Energy*, 85(9):2338–2348, 2011.

[27] Edgardo D Castronuovo and Joao A Pecas Lopes. Optimal operation and hydro storage sizing of a wind–hydro power plant. *International Journal of Electrical Power & Energy Systems*, 26(10):771–778, 2004.

[28] Sid Chi-Kin Chau, Jiajia Xu, Wilson Bow, and Khaled Elbassioni. Peer-to-peer energy sharing: Effective cost-sharing mechanisms and social efficiency. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pages 215–225. ACM, 2019.

[29] Min Chen and Gabriel A Rincon-Mora. Accurate electrical battery model capable of predicting runtime and iv performance. *IEEE transactions on energy conversion*, 21(2):504–511, 2006.

[30] SX Chen, Hoay Beng Gooi, and MingQiang Wang. Sizing of energy storage for microgrids. *IEEE Transactions on Smart Grid*, 3(1):142–151, 2012.

[31] Kangbeom Cheon, Jaehoon Kim, Moussa Hamadache, and Dongik Lee. On replacing pid controller with deep learning controller for dc motor system. *Journal of Automation and Control Engineering Vol*, 3(6), 2015.

[32] Jean-Michel Clairand, Mariano Arriaga, Claudio A Canizares, and Carlos Alvarez. Power generation planning of galapagos' microgrid considering electric vehicles and induction stoves. *IEEE Transactions on Sustainable Energy*, 2018.

[33] IBM ILOG CPLEX. V12. 1: Users manual for cplex. *International Business Machines Corporation*, 46(53):157, 2009.

[34] Indrajit Das, IEEE Member, and Claudio A Cañizares. Renewable energy integration in diesel-based microgrids at the canadian arctic. *Proceedings of the IEEE*, 2019.

[35] Jiaojiao Dong, Feng Gao, Xiaohong Guan, Qiaozhu Zhai, and Jiang Wu. Storage sizing with peak-shaving policy for wind farm based on cyclic markov chain model. *IEEE Transactions on Sustainable Energy*, 8(3):978–989, 2017.

[36] Bruce Dunn, Haresh Kamath, and Jean-Marie Tarascon. Electrical energy storage for the grid: a battery of choices. *Science*, 334(6058):928–935, 2011.

[37] Akram Eddahech, Olivier Briat, and J-M Vinassa. Neural networks based model and voltage control for lithium polymer batteries. In *Diagnostics for Electric Machines, Power Electronics & Drives (SDEMPED), 2011 IEEE International Symposium on*, pages 645–650. IEEE, 2011.

[38] EEMB. *LIR18650 cell*, 2016. Li-NMC cell specifications.

[39] Energy Regulation Quarterly. Electricity storage in north america. http://www.energyregulationquarterly.ca/articles/electricity-storage-in-north-america, 2019. Accessed: 2019-07-10.

[40] Enphase. Ac battery. https://enphase.com/en-us/products-and-services/storage, 2019. Accessed: 2019-05-10.

[41] Samira Fazlollahi, Pierre Mandel, Gwenaelle Becker, and Francois Maréchal. Methods for multi-objective investment and operating optimization of complex energy systems. *Energy*, 45(1):12–22, 2012.

[42] Summer Ferreira. Multiple uses in storage: Results of lithium-ion cells tested under stacked cycling profiles, 2015.

[43] Dariush Fooladivanda, Catherine Rosenberg, and Siddharth Garg. Energy storage and regulation: an analysis. *IEEE Transactions on Smart Grid*, 7(4):1813–1823, 2016.

[44] Philipp Fortenbacher, Johanna L Mathieu, and Göran Andersson. Modeling, identification, and optimal control of batteries for power system applications. In *2014 Power Systems Computation Conference*, pages 1–7. IEEE, 2014.

[45] Vincent François-Lavet, David Taralla, Damien Ernst, and Raphaël Fonteneau. Deep reinforcement learning solutions for energy microgrids management. In *European Workshop on Reinforcement Learning (EWRL 2016)*, 2016.

[46] Ran Fu, David Feldman, Robert Margolis, Mike Woodhouse, and Kristen Ardani. Us solar photovoltaic system cost benchmark: Q1 2017. Technical report, EERE Publication and Product Library, 2017.

[47] Jared L Gearhart, Kristin L Adair, Richard J Detry, Justin D Durfee, Katherine A Jones, and Nathaniel Martin. Comparison of open-source linear programming solvers. *Tech. Rep. SAND2013-8847*, 2013.

[48] Yashar Ghiassi-Farrokhfal, Fiodar Kazhamiaka, Catherine Rosenberg, and Srinivasan Keshav. Optimal design of solar pv farms with storage. *IEEE Transactions on Sustainable Energy*, 6(4):1586–1593, 2015.

[49] Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, and Catherine Rosenberg. Toward a realistic performance analysis of storage systems in smart grids. *IEEE Transactions on Smart Grid*, 6(1):402–410, 2015.

[50] Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, Catherine Rosenberg, and Florin Ciucu. Solar power shaping: An analytical approach. *IEEE Transactions on Sustainable Energy*, 6(1):162–170, 2015.

[51] Yashar Ghiassi-Farrokhfal, Catherine Rosenberg, Srinivasan Keshav, and Marie-Benedicte Adjaho. Joint optimal design and operation of hybrid energy storage systems. *IEEE Journal on Selected Areas in Communications*, 34(3):639–650, 2016.

[52] Paul Gilman, Nate Blair, Mark Mehos, Craig Christensen, Steve Janzou, and Chris Cameron. Solar advisor model user guide for version 2.0. *National Renewable Energy Laboratory, Golden, CO, Technical Report No. NREL/TP-670-43704*, 2008.

[53] Mevludin Glavic, Raphaël Fonteneau, and Damien Ernst. Reinforcement learning for electric power system decision and control: Past considerations and perspectives. *IFAC-PapersOnLine*, 50(1):6918–6927, 2017.

[54] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.

[55] P. P. Groumpos and G. Papageorgiou. An optimal sizing method for stand-alone photovoltaic power systems. *Solar Energy*, 38(5):341 – 351, Mar 1987.

[56] Omar Hafez and Kankar Bhattacharya. Optimal planning and design of a renewable energy based supply system for microgrids. *Renewable Energy*, 45:7–15, 2012.

[57] Krishnan S Hariharan and V Senthil Kumar. A nonlinear equivalent circuit model for lithium ion cells. *Journal of power sources*, 222:210–217, 2013.

[58] Pavithra Harsha and Munther Dahleh. Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy. *IEEE Transactions on Power Systems*, 30(3):1164–1181, 2015.

[59] Abubakar Sani Hassan, Liana Cipcigan, and Nick Jenkins. Optimal battery storage operation for pv systems with tariff incentives. *Applied Energy*, 203:422–441, 2017.

[60] Hongwen He, Rui Xiong, and Jinxin Fan. Evaluation of Lithium-ion battery equivalent circuit models for state of charge estimation by an experimental approach. *Energies*, 4(4):582–598, 2011.

[61] Ari Hentunen, Teemu Lehmuspelto, and Jussi Suomela. Time-domain parameter extraction method for thévenin-equivalent circuit battery models. *ieee transactions on energy conversion*, 29(3):558–566, 2014.

[62] Noriaki Hirose, Ryosuke Tajima, and Kazutoshi Sukigara. Mpc policy learning using dnn for human following control without collision. *Advanced Robotics*, 32(3):148–159, 2018.

[63] Xiaosong Hu, Shengbo Li, and Huei Peng. A comparative study of equivalent circuit models for li-ion batteries. *Journal of Power Sources*, 198:359–367, 2012.

[64] Tarun Huria, Massimo Ceraolo, Javier Gazzarri, and Robyn Jackey. High fidelity electrical model with thermal dependence for characterization and simulation of high power lithium battery cells. In *Electric Vehicle Conference (IEVC), 2012 IEEE International*, pages 1–8. IEEE, 2012.

[65] Robyn Jackey, Michael Saginaw, Pravesh Sanghvi, Javier Gazzarri, Tarun Huria, and Massimo Ceraolo. Battery model parameter estimation using a layered technique: An example using a lithium iron phosphate cell. Technical report, SAE Technical Paper, 2013.

[66] Milan Jain, Rachel K Kalaimani, Srinivasan Keshav, and Catherine Rosenberg. Using personal environmental comfort systems to mitigate the impact of occupancy prediction errors on hvac performance. *Energy Informatics*, 1(1):60, 2018.

[67] Jaehoon Jang and Jiyoon Yoo. Equivalent circuit evaluation method of lithium polymer battery using bode plot and numerical analysis. *IEEE Transactions on Energy Conversion*, 26(1):290–298, 2011.

[68] Yuming Jiang and Yong Liu. *Stochastic network calculus*, volume 1. Springer, 2008.

[69] Jiangliang Jin and Yunjian Xu. Optimal storage operation under demand charge. *IEEE Transactions on Power Systems*, 32(1):795–808, 2017.

[70] Rachel Kalaimani, Milan Jain, Srinivasan Keshav, and Catherine Rosenberg. On the interaction between personal comfort systems and centralized hvac systems in office buildings. *Advances in Building Energy Research*, pages 1–29, 2018.

[71] ND Kaushika, Nalin K Gautam, and Kshitiz Kaushik. Simulation model for sizing of stand-alone solar pv system with interconnected array. *Solar Energy Materials and Solar Cells*, 85(4):499–519, 2005.

[72] Hussein A Kazem, Tamer Khatib, and Kamaruzzaman Sopian. Sizing of a standalone photovoltaic/battery system at minimum cost for remote housing electrification in sohar, oman. *Energy and Buildings*, 61:108–115, 2013.

[73] Fiodar Kazhamiaka. Robust sizing. https://github.com/iss4e/Robust_Sizing, 2019.

[74] Fiodar Kazhamiaka, Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, and Catherine Rosenberg. Robust and practical approaches for solar pv and storage sizing. In *Proceedings of the Ninth International Conference on Future Energy Systems*, pages 146–156. ACM, 2018.

[75] Fiodar Kazhamiaka, Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, and Catherine Rosenberg. Comparison of different approaches for solar pv and storage sizing. In *revision with IEEE transactions on Sustainable Computing, invited paper*. IEEE, 2019.

[76] Fiodar Kazhamiaka, Patrick Jochem, Srinivasan Keshav, and Catherine Rosenberg. On the influence of jurisdiction on the profitability of residential photovoltaic-storage systems: A multi-national case study. *Energy Policy*, 109:428–440, 2017.

[77] Fiodar Kazhamiaka, Srinivasan Keshav, and Catherine Rosenberg. Adaptive battery control with neural networks. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pages 536–543. ACM, 2019.

[78] Fiodar Kazhamiaka, Srinivasan Keshav, Catherine Rosenberg, and Karl-Heinz Pettinger. Simple spec-based modeling of lithium-ion batteries. *IEEE Transactions on Energy Conversion*, 33(4):1757–1765, 2018.

[79] Fiodar Kazhamiaka, Catherine Rosenberg, and Srinivasan Keshav. Practical strategies for storage operation in energy systems: design and evaluation. *IEEE Transactions on Sustainable Energy*, 7(4):1602–1610, 2016.

[80] Fiodar Kazhamiaka, Catherine Rosenberg, and Srinivasan Keshav. Tractable lithium-ion storage models for optimizing energy systems. *Energy Informatics*, 2(1):4, 2019.

[81] Fiodar Kazhamiaka, Catherine Rosenberg, Srinivasan Keshav, and Karl-Heinz Pettinger. Li-ion storage models for energy system optimization: the accuracy-tractability tradeoff. In *Proceedings of the Seventh International Conference on Future Energy Systems*, page 17. ACM, 2016.

[82] W. D. Kellogg, M. H. Nehrir, G. Venkataramanan, and V. Gerez. Generation unit sizing and cost analysis for stand-alone wind, photovoltaic, and hybrid wind/PV systems. *IEEE Transactions on Energy Conversion*, 13(1):70 – 75, Mar 1998.

[83] David G Kendall. Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded markov chain. *The Annals of Mathematical Statistics*, pages 338–354, 1953.

[84] M Khalid and AV Savkin. A model predictive control approach to the problem of wind power smoothing with controlled battery storage. *Renewable Energy*, 35(7):1520–1526, 2010.

[85] Rajab Khalilpour and Anthony Vassallo. Planning and operation scheduling of pv-battery systems: A novel methodology. *Renewable and Sustainable Energy Reviews*, 53:194–208, 2016.

[86] Arash Khatamianfar, Muhammad Khalid, Andrey V Savkin, and Vassilios G Agelidis. Improving wind farm dispatch in the australian electricity market with battery energy storage using model predictive control. *IEEE Transactions on Sustainable Energy*, 4(3):745–755, 2013.

[87] Tamer Khatib, Ibrahim A Ibrahim, and Azah Mohamed. A review on sizing methodologies of photovoltaic array and storage battery in a standalone photovoltaic system. *Energy Conversion and Management*, 120:430–448, 2016.

[88] T. Khatiba, A. Mohameda, and K. Sopian. Renewable and sustainable energy reviews. *Renewable Energy*, 22(1):454 – 465, Jun 2013.

[89] Taesic Kim and Wei Qiao. A hybrid battery model capable of capturing dynamic circuit characteristics and nonlinear capacity effects. *IEEE Transactions on Energy Conversion*, 26(4):1172–1180, 2011.

[90] Reinhardt Klein, Nalin A Chaturvedi, Jake Christensen, Jasim Ahmed, Rolf Findeisen, and Aleksandar Kojic. Electrochemical model based observer design for a lithium-ion battery. *IEEE Transactions on Control Systems Technology*, 21(2):289–301, 2013.

[91] MC Knauff, Chris J Dafis, Dagmar Niebur, Harry G Kwatny, CO Nwankpa, and J Metzer. Simulink model for hybrid power system test-bed. In *Electric Ship Technologies Symposium, 2007. ESTS'07. IEEE*, pages 421–427. IEEE, 2007.

[92] Kokam. *Li-ion Polymer Cell*, 2016. Li-NMC cell specifications.

[93] Michael Koller, Theodor Borsche, Andreas Ulbig, and Göran Andersson. Defining a degradation cost function for optimal control of a battery energy storage system. In *2013 IEEE Grenoble Conference*, pages 1–6. IEEE, 2013.

[94] F. Kong, C. Dong, X. Liu, and H. Zeng. Blowing hard is not all we want: Quantity vs quality of wind power in the smart grid. In *IEEE INFOCOM*, pages 2813 – 2821, April 2014.

[95] L Kools and F Phillipson. Data granularity and the optimal planning of distributed generation. *Energy*, 112:342–352, 2016.

[96] A. Kornelakis. Multiobjective particle swarm optimization for the optimal design of photovoltaic grid-connected systems. *IET Renewable Power Generation*, 84(12):2022 – 2033, December 2010.

[97] A Kornelakis and E. Koutroulis. Methodology for the design optimisation and the economic analysis of grid-connected photovoltaic systems. *IET Renewable Power Generation*, 3(4):476 – 492, December 2009.

[98] Magnus Korpaas, Arne T Holen, and Ragne Hildrum. Operation and sizing of energy storage for wind power plants in a market system. *International Journal of Electrical Power & Energy Systems*, 25(8):599–606, 2003.

[99] E. Koutroulis, D. Kolokotsab, and A. Potirakisa K. Kalaitzakis. Methodology for optimal sizing of stand-alone photovoltaic/wind-generator systems using genetic algorithms. *Solar Energy*, 80(9):1072 – 1088, Sep 2006.

[100] Pratyush Kumar. A network calculus foundation for smart-grids where demand and supply vary in space and time. In *Proceedings of the Seventh International Conference on Future Energy Systems*, e-Energy '16, pages 1–11, New York, NY, USA, 2016. ACM.

[101] Long Lam, Pavol Bauer, and Erik Kelder. A practical circuit-based model for li-ion battery cells in electric vehicle applications. In *2011 IEEE 33rd International Telecommunications Energy Conference (INTELEC)*, pages 1–9. IEEE, 2011.

[102] Andrzej Lebkowski. Temperature, overcharge and short-circuit studies of batteries used in electric vehicles. *Przeglad Elektrotechniczny*, 93(5):67–73, 2017.

[103] Leclanché. *LecCell 30Ah High Energy*, 02 2014. Lithium-Titanate cell specifications.

[104] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[105] Nian Liu, Xinghuo Yu, Cheng Wang, Chaojie Li, Li Ma, and Jinyong Lei. Energy-sharing model with price-based demand response for microgrids of peer-to-peer prosumers. *IEEE Transactions on Power Systems*, 32(5):3569–3583, 2017.

[106] Stefano Longo, Eric C Kerrigan, Keck Voon Ling, and George A Constantinides. A parallel formulation for predictive control with nonuniform hold constraints. *Annual Reviews in Control*, 35(2):207–214, 2011.

[107] Juan Miguel Gonzalez Lopez, Edris Pouresmaeil, Claudio A Canizares, Kankar Bhattacharya, Abolfazl Mosaddegh, and Bharatkumar V Solanki. Smart residential load simulator for energy management in smart grids. *IEEE Transactions on Industrial Electronics*, 66(2):1443–1452, 2018.

[108] Srdjan M Lukic, Jian Cao, Ramesh C Bansal, Fernando Rodriguez, and Ali Emadi. Energy storage systems for automotive applications. *IEEE Transactions on industrial electronics*, 55(6):2258–2267, 2008.

[109] Johannes Mayer. Electricity production and spot prices in germany 2014. *Report from Fraunhofer Institute for Solar Energy Systems, Germany*, 2014.

[110] Brida Mbuwir, Frederik Ruelens, Fred Spiessens, and Geert Deconinck. Battery energy management in a microgrid using batch reinforcement learning. *Energies*, 10(11):1846, 2017.

[111] Mahmoud Mehrabankhomartash, Mohammad Rayati, Aras Sheikhi, and Ali Mohammad Ranjbar. Practical battery size optimization of a pv system by considering individual customer damage function. *Renewable and Sustainable Energy Reviews*, 67:36–50, 2017.

[112] A. Mellita, S. A. Kalogiroub, L. Hontoriac, and S. Shaari. Artificial intelligence techniques for sizing photovoltaic systems: A review. *Renewable and Sustainable Energy Reviews*, 13(2):406 – 419, Feb 2009.

[113] Alan Millner. Modeling lithium ion battery degradation in electric vehicles. In *Innovative Technologies for an Efficient and Reliable Electricity Supply (CITRES), 2010 IEEE Conference on*, pages 349–356. IEEE, 2010.

[114] Joydeep Mitra. Reliability-based sizing of backup storage. *IEEE Transactions on Power Systems*, 25(2):1198–1199, 2010.

[115] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[116] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[117] Jorge Moreno, Micah E Ortúzar, and Juan W Dixon. Energy-management system for a hybrid electric vehicle, using ultracapacitors and neural networks. *IEEE transactions on Industrial Electronics*, 53(2):614–623, 2006.

[118] Thomas Morstyn, Andrey V Savkin, Branislav Hredzak, and Vassilios G Agelidis. Distributed sliding mode control for multi-module battery energy storage system state of charge balancing. In *2016 IEEE Conference on Control Applications (CCA)*, pages 47–51. IEEE, 2016.

[119] John M Mulvey, Robert J Vanderbei, and Stavros A Zenios. Robust optimization of large-scale systems. *Operations research*, 43(2):264–281, 1995.

[120] S Nejad, DT Gladwin, and DA Stone. A systematic review of lumped-parameter equivalent circuit models for real-time estimation of lithium-ion battery states. *Journal of Power Sources*, 316:183–196, 2016.

[121] M. H. Nehrir Nelson, D. B. and C. Wang. Unit sizing and cost analysis of stand-alone hybrid wind/PV/fuel cell power generation systems. *Renewable Energy*, 31(10):1641 – 1656, Aug 2006.

[122] P. G. Nikhil and D. Subhakar. An improved algorithm for photovoltaic system sizing. *Energy Procedia*, 14:1134 – 1142, 2012.

[123] Seyyed Mostafa Nosratabadi, Rahmat-Allah Hooshmand, and Eskandar Gholipour. A comprehensive review on microgrid and virtual power plant concepts employed for distributed energy resources scheduling in power systems. *Renewable and Sustainable Energy Reviews*, 67:341–363, 2017.

[124] Björn Nykvist and Måns Nilsson. Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5(4):329–332, 2015.

[125] Chiemeka Onyeka Okoye and Oğuz Solyalı. Optimal sizing of stand-alone photovoltaic systems in residential buildings. *Energy*, 126:573–584, 2017.

[126] Daniel E Olivares, Jose D Lara, Claudio A Cañizares, and Mehrdad Kazerani. Stochastic-predictive energy management system for isolated microgrids. *IEEE Transactions on Smart Grid*, 6(6):2681–2693, 2015.

[127] Noshin Omar, Peter Van den Bossche, Thierry Coosemans, and Joeri Van Mierlo. Peukert revisitedcritical appraisal and need for modification for lithium-ion batteries. *Energies*, 6(11):5625–5641, 2013.

[128] Keiron O'Shea and Ryan Nash. An introduction to convolutional neural networks. *arXiv preprint arXiv:1511.08458*, 2015.

[129] Minggao Ouyang, Xuning Feng, Xuebing Han, Languang Lu, Zhe Li, and Xiangming He. A dynamic capacity degradation model and its applications considering varying load for a large format li-ion battery. *Applied Energy*, 165:48–59, 2016.

[130] Quan Ouyang, Jian Chen, Jian Zheng, and Huazhen Fang. Optimal cell-to-cell balancing topology design for serially connected lithium-ion battery packs. *IEEE Transactions on Sustainable Energy*, 9(1):350–360, 2017.

[131] M Rosa Palacín and Anne de Guibert. Why do batteries fail? *Science*, 351(6273):1253292, 2016.

[132] Pablo Pedregal. *Introduction to optimization*, volume 46. Springer Science & Business Media, 2006.

[133] Changhao Piao, Xiaoyong Yang, Cong Teng, and HuiQian Yang. An improved model based on artificial neural networks and thevenin model for nickel metal hydride power battery. In *Optics Photonics and Energy Engineering (OPEE), 2010 International Conference on*, volume 1, pages 115–118. IEEE, 2010.

[134] Gregory L Plett. Extended kalman filtering for battery management systems of lipb-based hev battery packs: Part 2. modeling and identification. *Journal of power sources*, 134(2):262–276, 2004.

[135] Baihong Qi, Mohammad Rashedi, and Omid Ardakanian. Energyboost: Learning-based control of home batteries. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, e-Energy '19, pages 239–250, New York, NY, USA, 2019. ACM.

[136] S Joe Qin and Thomas A Badgwell. A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764, 2003.

[137] Majid Raeis, Almut Burchard, and Jörg Liebeherr. Analysis of the leakage queue: A queuing model for energy storage systems with self-discharge. *CoRR*, abs/1710.09506, 2017.

[138] J Anthony Rossiter. *Model-based predictive control: a practical approach*. CRC press, 2017.

[139] Y. Ru, J. Kleissl, and S. Martinez. Storage size determination for grid-connected photovoltaic systems. *IEEE Transactions on Sustainable Energy*, 4(1):68 – 81, Jan 2013.

[140] Tim Salsbury, Prashant Mhaskar, and S Joe Qin. Predictive control methods to improve energy efficiency and reduce demand in buildings. *Computers & Chemical Engineering*, 51:77–85, 2013.

[141] John G Saw, Mark CK Yang, and Tse Chin Mo. Chebyshev inequality with estimated mean and variance. *The American Statistician*, 38(2):130–132, 1984.

[142] Riccardo Scattolini and Patrizio Colaneri. Hierarchical model predictive control. In *2007 46th IEEE Conference on Decision and Control*, pages 4803–4808. IEEE, 2007.

[143] Aden Seaman, Thanh-Son Dao, and John McPhee. A survey of mathematics-based equivalent-circuit and electrochemical battery models for hybrid and electric vehicle simulation. *Journal of Power Sources*, 256:410–423, 2014.

[144] N. Sharma and V. Siddhartha. Stochastic techniques used for optimization in solar systems: A review. *Renewable and Sustainable Energy Reviews*, 16(3):1399 – 1411, Apr 2012.

[145] V. K. Sharma, A. Colangelo, and G. Spagna. Photovoltaic technology: Basic concepts, sizing of a stand alone photovoltaic system for domestic applications and preliminary economic analysis. *Energy Conversion and Management*, 36(3):161 – 174, Mar 1995.

[146] Clarence M Shepherd. Design of primary and secondary cells ii. an equation describing battery discharge. *Journal of the Electrochemical Society*, 112(7):657–664, 1965.

[147] S. Singla, Y. Ghiassi-Farrokhfal, and S. Keshav. Using storage to minimize carbon footprint of diesel generators for unreliable grids. *IEEE Transactions on Sustainable Energy*, 5(4):1270–1277, Oct 2014.

[148] Kandler A Smith, Christopher D Rahn, and Chao-Yang Wang. Control oriented 1d electrochemical model of lithium ion battery. *Energy Conversion and management*, 48(9):2565–2578, 2007.

[149] Catherine Soanes and Angus Stevenson. *The oxford english reference dictionary, 2nd edition*. Oxford University Press, 2003.

[150] Bharatkumar V Solanki, Akash Raghurajan, Kankar Bhattacharya, and Claudio A Cañizares. Including smart loads for optimal demand response in integrated energy management systems for isolated microgrids. *IEEE Transactions on Smart Grid*, 8(4):1739–1748, 2015.

[151] C. Soras and V. Makios. A novel method for determining the optimum size of standalone photovoltaic systems. *Solar Cells*, 25(2):127 – 142, Mar 1987.

[152] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[153] Bartolomeo Stellato, Bart PG Van Parys, and Paul J Goulart. Multivariate chebyshev inequality with estimated mean and variance. *The American Statistician*, 71(2):123–127, 2017.

[154] stem.com. Ai powered energy storage, March 2019.

[155] Rudi Studer, V Richard Benjamins, and Dieter Fensel. Knowledge engineering: principles and methods. *Data & knowledge engineering*, 25(1):161–197, 1998.

[156] Sun Sun, Fiodar Kazhamiaka, Srinivasan Keshav, and Catherine Rosenberg. Using synthetic traces for robust energy system sizing. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pages 251–262. ACM, 2019.

[157] Chee Wei Tan, Tim C Green, and Carlos A Hernandez-Aramburo. A stochastic method for battery sizing with uninterruptible-power and demand shift capabilities in pv (photovoltaic) systems. *Energy*, 35(12):5082–5092, 2010.

[158] Mike Taylor, Joyce McLaren, Karlynn Cory, Ted Davidovich, John Sterling, and Miriam Makhyoun. Value of solar. program design and implementation considerations. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2015.

[159] Sercan Teleke, Mesut E Baran, Subhashish Bhattacharya, and Alex Q Huang. Optimal control of battery energy storage for wind farm dispatching. *IEEE Transactions on Energy Conversion*, 25(3):787–794, 2010.

[160] Sercan Teleke, Mesut E Baran, Subhashish Bhattacharya, and Alex Q Huang. Rule-based control of battery energy storage for dispatching intermittent renewable sources. *Sustainable Energy, IEEE Transactions on*, 1(3):117–124, 2010.

[161] Tesla corp. Powerwall 2. https://www.tesla.com/powerwall, 2019. Accessed: 2019-05-10.

[162] Kannan Thirugnanam, Ezhil Reena Joy TP, Mukesh Singh, and Praveen Kumar. Mathematical modeling of li-ion battery using genetic algorithm approach for v2g applications. *IEEE transactions on Energy conversion*, 29(2):332–343, 2014.

[163] Michael James Tippett, Chee Keong Tan, and Jie Bao. Non-constant prediction-step mpc for processes with multi-scale dynamics. *IFAC Proceedings Volumes*, 47(3):3068–3073, 2014.

[164] Nikita Tomin, Victor Kurbatsky, and Michael Negnevitsky. The concept of the deep learning-based system" artificial dispatcher" to power system control and dispatch. *arXiv preprint arXiv:1805.05408*, 2018.

[165] Olivier Tremblay and Louis-A Dessaint. Experimental validation of a battery dynamic model for ev applications. *World Electric Vehicle Journal*, 3(1):1–10, 2009.

[166] Olivier Tremblay, Louis-A Dessaint, and Abdel-Illah Dekkiche. A generic battery model for the dynamic simulation of hybrid electric vehicles. In *Vehicle Power and Propulsion Conference, 2007. VPPC 2007. IEEE*, pages 284–289. IEEE, 2007.

[167] Pravin M Vaidya. Speeding-up linear programming using fast matrix multiplication. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 332–337. IEEE, 1989.

[168] Alexandra Von Meier. *Electric power systems: a conceptual introduction*. John Wiley & Sons, 2006.

[169] Caisheng Wang and M Hashem Nehrir. Power management of a stand-alone wind/photovoltaic/fuel cell energy system. *IEEE transactions on energy conversion*, 23(3):957–967, 2008.

[170] Fu-Kwun Wang and Tadele Mamo. A hybrid model based on support vector regression and differential evolution for remaining useful lifetime prediction of lithium-ion batteries. *Journal of Power Sources*, 401:49–54, 2018.

[171] K. Wang, F. Ciucu, C. Lin, and S. H. Low. A stochastic power network calculus for integrating renewable energy sources into the power grid. *IEEE Journal on Selected Areas in Communications*, 30(6):1037–1048, July 2012.

[172] Kai Wang, Florin Ciucu, Chuang Lin, and Steven H Low. A stochastic power network calculus for integrating renewable energy sources into the power grid. *IEEE Journal on Selected Areas in Communications*, 30(6):1037–1048, 2012.

[173] Le Yi Wang, Caisheng Wang, George Yin, Feng Lin, Michael P Polis, Caiping Zhang, and Jiuchun Jiang. Balanced control strategies for interconnected heterogeneous battery systems. *IEEE Transactions on Sustainable Energy*, 7(1):189–199, 2016.

[174] Johannes Weniger, Tjarko Tjaden, and Volker Quaschning. Sizing of residential pv battery systems. *Energy Procedia*, 46:78–87, 2014.

[175] Herman LN Wiegman. *Battery state estimation and control for power buffering applications*. University of Wisconsin–Madison, 1999.

[176] WorldInfo. Energy consumption in canada. https://www.worlddata.info/america/canada/energy-consumption.php, 2019. Accessed: 2019-07-10.

[177] xcelenergy.com. Time-of-day pricing, March 2019.

[178] Le Xie, Yingzhong Gu, Ali Eskandari, and Mehrdad Ehsani. Fast mpc-based coordination of wind power and battery energy storage systems. *Journal of Energy Engineering*, 138(2):43–53, 2012.

[179] Bolun Xu, Alexandre Oudalov, Andreas Ulbig, Göran Andersson, and Daniel S Kirschen. Modeling of lithium-ion battery degradation for cell life assessment. *IEEE Transactions on Smart Grid*, 9(2):1131–1140, 2016.

[180] Tianfang Xu, Albert J Valocchi, Jaesik Choi, and Eyal Amir. Use of machine learning methods to reduce predictive error of groundwater models. *Groundwater*, 52(3):448–460, 2014.

[181] Hongxing Yang, Wei Zhou, Lin Lu, and Zhaohong Fang. Optimal sizing method for stand-alone hybrid solar–wind system with lpsp technology by using genetic algorithm. *Solar energy*, 82(4):354–367, 2008.

[182] Low Wen Yao, JA Aziz, Pui Yee Kong, and NRN Idris. Modeling of lithium-ion battery using matlab/simulink. In *Industrial Electronics Society, IECON 2013-39th Annual Conference of the IEEE*, pages 1729–1734. IEEE, 2013.

[183] Yonghuang Ye, Yixiang Shi, Ningsheng Cai, Jianjun Lee, and Xiangming He. Electro-thermal modeling and experimental validation for lithium ion battery. *Journal of Power Sources*, 199:227–238, 2012.

[184] XI Yu-Geng, LI De-Wei, and Lin Shu. Model predictive control - status and challenges. *Acta Automatica Sinica*, 39(3):222–236, 2013.

[185] Sheng Shui Zhang. The effect of the charging protocol on the cycle life of a Li-ion battery. *Journal of power sources*, 161(2):1385–1391, 2006.

[186] Tianhao Zhang, Gregory Kahn, Sergey Levine, and Pieter Abbeel. Learning deep control policies for autonomous aerial vehicles with mpc-guided policy search. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 528–535. IEEE, 2016.

[187] Yan Zhang, Fanlin Meng, Rui Wang, Wanlu Zhu, and Xiao-Jun Zeng. A stochastic mpc based approach to integrated energy management in microgrids. *Sustainable cities and society*, 41:349–362, 2018.

[188] Dong Zhou, Long Xue, Yijia Song, and Jiayu Chen. On-line remaining useful life prediction of lithium-ion batteries based on the optimized gray model gm (1, 1). *Batteries*, 3(3):21, 2017.

[189] Ting Zhu, Aditya Mishra, David Irwin, Navin Sharma, Prashant Shenoy, and Don Towsley. The case for efficient renewable energy management in smart homes. In *Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 67–72. ACM, 2011.